Simulating from the Heston Model: A Gamma Approximation Scheme

Submitted to Monte Carlo Methods and Applications

Jean-François Bégin[1], Mylène Bédard[2], and Patrice Gaillardetz*[3]

[1]Department of Decision Sciences, HEC Montréal
[2]Department of Mathematics and Statistics, Université de Montréal
[3]Department of Mathematics and Statistics, Concordia University

First draft: August 10, 2012
This version: June 3, 2014

Abstract

The Heston model is appealing as it possesses a stochastic volatility term as well as semi-closed formulas for pricing European options. Unfortunately, few simulation schemes for this model can handle the violation of the Feller Condition.

An algorithm based on the exact scheme of Broadie and Kaya to simulate price paths under the Heston model is introduced. In order to increase the speed of their exact method, we use a gamma approximation. According to Stewart et al., it is possible to approximate a complex gamma convolution (similar to the representation given by Glasserman and Kim) by a simple moment-matched gamma distribution.

We also perform a review of popular simulation schemes for the Heston model and validate our approach through a simulation study. The gamma approximation scheme appears to yield small biases on European and Asian option prices when compared to the most popular schemes.

Keywords: Stochastic volatility, Heston model, Simulation schemes, Gamma expansion, Asian options.

1 Introduction

Financial stocks are often modelled by stochastic differential equations (SDEs). These equations describe the behaviour of the underlying asset and also of certain model parameters. Nowadays, the models retaining our attention have stochastic volatility. They are known to allow for a

*The authors acknowledge the financial support from the National Science and Engineering Research Council of Canada (NSERC).

better calibration to market data; they also efficiently capture the smile of the volatility observed in financial securities.

Until recently, the Black-Scholes-Merton (BSM) model [5, 23] was widely used. For most contingent claims, semi-closed formulas exist for this model, which makes it very attractive from a practical point of view. However, the BSM model includes very coarse assumptions such as a constant volatility and a deterministic asset growth rate. These shortfalls, combined to several financial crashes and the introduction of complex products, have forced financial analysts to develop new models. Heston [16] proposes a model based on the square root process with mean reversion to express variance. This model became very popular among practitioners as it admits semi-closed formulas for pricing European call options. Moreover, the variance process (square root process) is widely applied in finance due to the availability of several analytic results about this SDE; for instance, see the short rate model of Cox, Ingersoll and Ross [8].

Although there exists formulas for pricing European call options under the Heston model, there does not exist such expressions for complex products involving path dependency. To price such products, we thus rely on Monte Carlo simulation techniques. Despite the fact that the Heston model is nearly twenty years old, efficient simulation procedures have interested only a handful of individuals.

Generally, the scheme named after Leonhard Euler and Gisiro Maruyama and the one introduced by Milstein [25] are the most efficient. They are also the easiest to implement: these methods can be used with virtually any SDE. However, under the Heston model, these techniques do not work very well when the time steps are long or when the Feller Condition is not satisfied.

In an attempt to overcome these drawbacks, several researchers have modified these methods for use with the Heston model. Lord, Koekkoek and van Dijk [21] consider a large number of fixes for the Euler-Maruyama method. The most convincing, called full truncation, is designed to minimize the bias on European call option prices. Kahl and Jäckel [19] suggest discretizing the variance process $\{V(t) : t \geq 0\}$ using an implicit Milstein scheme, coupled with their own discretization method for the asset process. Unfortunately, according to Andersen [2], this scheme is significantly biased when the Feller Condition is violated.

On a different note, Broadie and Kaya [7] introduce a method said exact to simulate from the Heston model. Essentially, by obtaining a value for the variance $V(t)$ at some specific time t, one can then use Fourier inversion techniques to generate the conditional integrated variance over time given the bounds of the integral. In their paper, they derive a closed-form expression for the characteristic function of this integral, which renders its simulation possible; subsequently, one can easily recover the price. This method, though elegant, have serious limitations: the computational effort required for implementing this algorithm makes it practically unusable (see van Haastrecht and Pelsser [31]). Nonetheless, the theory of [7] generated enthusiasm and several researchers used this idea to develop their own technique.

Among these, Smith [28] proposes an almost exact algorithm by substituting the arithmetic and geometric averages in the characteristic function by a weighted average of these two quantities. This results in a characteristic function depending on two variables instead of three. For fixed values of these two variables, corresponding values of the characteristic function are then cached in order to accelerate the simulation algorithm. However, this modification barely accelerates the exact method of Broadie and Kaya [7]: Smith's scheme is about 70 times slower than the best methods, according to van Haastrecht and Pelsser [31].

Glasserman and Kim [15] find that the integrated variance over time, available given the bounds of the integral, can be explicitly represented by a sum of infinite mixtures of gamma random variables. This random variable can be easily generated by truncating the infinite series. Consequently, the longest step of Broadie and Kaya's [7] method becomes much faster.

Andersen [2] takes a similar avenue; he presents an approach aiming at efficiently approximating the variance using moment-matching methods. The idea behind his scheme, called quadratic exponential (QE hereafter), is to match the moments of a Gaussian density where the probability under zero is assigned to a Dirac delta function at the origin. This reproduces the asymptotic behaviour of the variance process for large values of $V(t)$. For smaller values of $V(t)$, an alternative density based on an exponential and a Dirac delta function is preferred. This algorithm is considered by many, including Tse and Wan [30], as one of the best algorithms for simulating from the Heston model.

Instead of working with the variance SDE, Zhu [32] works with the root of the variance SDE, also called the volatility process. Using some approximations for the new parameters of the process along with the Euler-Maruyama discretization, he shows that it is possible to obtain an algorithm that gives very similar results to Andersen [2]. However, this scheme does not work well when the Feller Condition is violated.

Finally, Tse and Wan [30] consider a biased approximation based on the inverse Gaussian distribution. They show that the integrated variance of the Heston model converges to an inverse Gaussian when the time step goes to infinity. Using this result, the authors develop an algorithm that uses an efficient approximation of the inverse Gaussian to simulate from the Heston model.

The scheme introduced in this paper relies on the basics of Broadie and Kaya; in order to speed up the latter scheme, we use a gamma distribution to approximate the integrated variance over time. By caching the moments of the real integrated variance over time, we manage to efficiently approximate this random variable. The contribution of this paper is two-fold: we first propose a sampling algorithm for the Heston model, and then we compare our method to other schemes used in practice. The simulation studies are based on realistic and challenging parameters (similar to those observed in market data). Our applications focus on implied model parameters that can be considered quite extreme.

The ability to generate observations for the integrated variance over time is crucial when pricing path-dependent securities. The gamma approximation algorithm introduced in this paper allows us to price any kind of path-dependent derivative. The simulation of the integrated variance over time is also useful in other contexts. One might be interested, for instance, in computing the greeks for some derivatives. These quantities need to be computed through simulations, and the integrated variance over time is a quantity that is recurrently needed in such calculations. Hence, the proposed sampling scheme could be used to accelerate the valuation of greeks. The implementation of the conditional expectation method (see [17] for an example and [6] for an application to the greeks under the Heston framework) often requires generating observations from the distribution of the integrated variance over time, and so constitutes another example of the applicability of the proposed algorithm.

This paper is organized as follows: in Section 2, we present the Heston dynamics and we describe important results about the processes involved. Section 3 is devoted to a complete analysis of popular simulation schemes for the Heston model. In Section 4, we introduce a new sampling scheme based on a gamma approximation. Section 5 compares our approach with those mentioned earlier. Section 6 concludes.

2 Heston Framework

The Heston dynamics are presented; the moments, characteristic function, and probability distribution of the integrated variance over time are then introduced.

2.1 Heston dynamics

The Heston model is defined, under the risk-neutral probability measure, by two coupled SDEs: one for the asset price and another for the variance. Let $S = \{S(t); t \geq 0\}$ be the asset price process and $V = \{V(t); t \geq 0\}$ the variance process.

Proposition 1 (Heston model)**.** *Under the risk-neutral probability measure $\mathbb{Q}$, the Heston model is given by*

$$\begin{cases} dS(t) = rS(t)dt + \sqrt{V(t)}S(t)dW_S(t) & (1) \\ dV(t) = \kappa(\theta - V(t))\,dt + \sigma\sqrt{V(t)}dW_V(t) & (2) \\ S(0) \text{ and } V(0) \text{ are deterministic initial conditions of the SDEs} \end{cases}$$

where

- *r is the risk-free rate,*

- κ *is the speed of the mean-reversion,*
- θ *is the average level of the variance process* V*,*
- σ *is the volatility of the variance process* V*,*
- $W_S(t)$ *and* $W_V(t)$ *are the time-*t *values of two Brownian motions and*
- ρ *is the correlation coefficient between the Brownian motions.*

We refer the reader to Moodley [26] for a derivation of the model under the risk-neutral measure. From a computational viewpoint, it is generally convenient to apply a transformation to the asset price process S, which is similar to a geometric Brownian motion. To this end, let $X = \log(S)$; applying Itô's Lemma to (1) and (2) leads to the following SDEs.

Proposition 2 (Heston model, log returns dynamics)**.** *Under the risk-neutral probability measure* $\mathbb{Q}$*, the Heston model is given by*

$$\begin{cases} dX(t) = \left(r - \frac{1}{2}V(t)\right)dt + \sqrt{V(t)}dW_X(t) & (3) \\ dV(t) = \kappa(\theta - V(t))\,dt + \sigma\sqrt{V(t)}dW_V(t) & (4) \\ X(0) \text{ and } V(0) \text{ are deterministic initial conditions of the SDEs} \end{cases}$$

where $X(t) = \log(S(t))$*,* $W_X(t) = W_S(t)$*, and the parameters are as in Definition 1.*

Proof. See [2]. □

The performance of simulation schemes sometimes depends on the satisfaction of a specific condition, known as the Feller Condition [11], for the square root process.

Proposition 3 (Feller Condition)**.** *Assume that* $V(0) > 0$*. If* $2\kappa\theta \geq \sigma^2$*, then the variance process* V *cannot reach zero.*

A proof of this proposition can be found in [2]. Note that when the Feller Condition is violated (i.e. $2\kappa\theta < \sigma^2$) then the origin is accessible and strongly reflecting; in such cases, the likelihood of hitting zero is often quite significant. In Figure 1, the cumulative distribution functions (cdf) of $V(T)$ given $V(t)$ for two different sets of parameters illustrate the Feller Condition. One can easily see that the cumulative distribution function is non-zero near the origin if the Feller Condition is violated (dashed line); this is however not the case when the condition is satisfied (solid line).

2.2 Analysis of the variance dynamics

The mean-reverting square root process V is identical to the process used by Cox et al. [8] to model short interest rates. Mean-reversion is a realistic (and thus desired) property for a variance

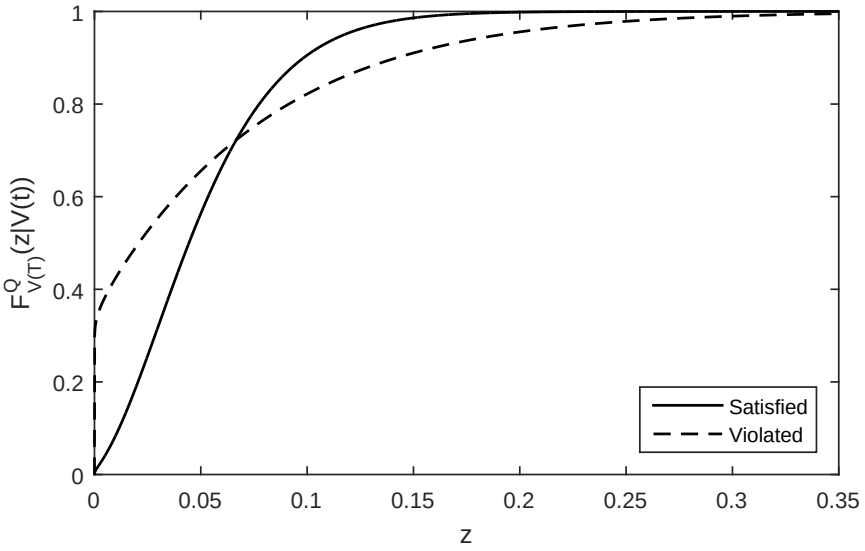

Figure 1: Plot of $F^{\mathbb{Q}}_{V(T)}(z|V(t))$ against z for two specific cases : one where the Feller Condition is satisfied (solid line) and another where it is violated (dashed line).

process: broadly speaking, it means that, over time, the process tends to move to its long run average. This process is well-documented and several related analytical results are available; we refer the reader to [3], [8] and [10] for more details on the mathematical aspects of this SDE. We now list two important results, which are proven in [2].

Proposition 4 (Cumulative distribution function of $V(T)$ given $V(t)$)**.** *Let $F_{\chi^2}(z;\nu,\lambda)$ be the cdf of the non-central chi-square distribution with non-centrality parameter λ and ν degrees of freedom,*

$$F_{\chi^2}(z;\nu,\lambda) = e^{-\frac{\lambda}{2}} \sum_{j=0}^{\infty} \frac{\left(\frac{\lambda}{2}\right)^j}{j!2^{\frac{\nu}{2}+j}\Gamma\left(\frac{\nu}{2}+j\right)} \int_0^z x^{\frac{\nu}{2}+j-1} e^{-\frac{x}{2}}\, dx. \tag{5}$$

Let

$$d = \frac{4\kappa\theta}{\sigma^2} \tag{6}$$

and

$$n(t,T) = \frac{4\kappa e^{-\kappa(T-t)}}{\sigma^2\left(1-e^{-\kappa(T-t)}\right)}, \quad T > t,$$

where κ, θ and σ are as in Proposition 1. Conditional on $V(t)$, the cdf of $V(T)$ is

$$F^{\mathbb{Q}}_{V(T)}(z|V(t)) = F_{\chi^2}\left(\frac{zn(t,T)}{e^{-\kappa(T-t)}}; d, V(t)n(t,T)\right).$$

Closed-form solutions are available for the moments of the non-central chi-square distribution; expressions for the expectation and variance of $V(T)$ conditional on $V(t)$ are provided below.

Proposition 5 (Moments of $V(T)$ given $V(t)$)**.** *Conditional on $V(t)$, the first two moments of $V(T)$*

are expressed as

$$\mathbb{E}[V(T)|V(t)] = \theta + (V(t) - \theta)e^{-\kappa(T-t)} \tag{7}$$

and

$$\mathrm{Var}[V(T)|V(t)] = \frac{V(t)\sigma^2 e^{-\kappa(T-t)}}{\kappa}\left(1 - e^{-\kappa(T-t)}\right) + \frac{\theta\sigma^2}{2\kappa}\left(1 - e^{-\kappa(T-t)}\right)^2, \tag{8}$$

where κ, θ and σ are as in Proposition 1.

2.3 Cumulative distribution function of the integrated variance over time

The integrated variance over time (quadratic variation increment) under Heston, defined as

$$IV(t,T) = \int_t^T V(s)\, ds \tag{9}$$

has been studied by Broadie and Kaya [7]; in particular, the authors derive an expression for the cdf of this random variable, which is now stated. Hereafter, the parameters κ, θ and σ refer to the parameters defined in Proposition 1.

Proposition 6. (Cumulative distribution function of the integrated variance over time). *Given $V(t)$ and $V(T)$, the cdf of the integrated variance over time under the Heston model is expressed as*

$$F^{\mathbb{Q}}_{IV(t,T)}(z|V(t),V(T)) = \frac{2}{\pi}\int_0^\infty \frac{\sin(\phi z)}{\phi}\mathrm{Re}\left(\Xi(\phi)\right)\, d\phi,$$

where

$$\begin{aligned}\Xi(\phi) = {} & \frac{\gamma(\phi)e^{-\frac{1}{2}(\gamma(\phi)-\kappa)(T-t)}[1-e^{-\kappa(T-t)}]}{\kappa[1-e^{-\gamma(\phi)(T-t)}]} \\ & \times \exp\left\{\frac{V(t)+V(T)}{\sigma^2}\left(\frac{\kappa[1+e^{-\kappa(T-t)}]}{1-e^{-\kappa(T-t)}} - \frac{\gamma(\phi)[1+e^{-\gamma(\phi)(T-t)}]}{1-e^{-\gamma(\phi)(T-t)}}\right)\right\} \\ & \times \mathcal{I}_{\frac{d}{2}-1}\left(\sqrt{V(t)V(T)}\frac{4\gamma(\phi)e^{-\frac{1}{2}\gamma(\phi)(T-t)}}{\sigma^2[1-e^{-\gamma(\phi)(T-t)}]}\right)\Bigg/ \\ & \mathcal{I}_{\frac{d}{2}-1}\left(\sqrt{V(t)V(T)}\frac{4\kappa e^{-\frac{1}{2}\kappa(T-t)}}{\sigma^2[1-e^{-\kappa(T-t)}]}\right),\end{aligned} \tag{10}$$

with

$$\gamma(\phi) = \sqrt{\kappa^2 - 2\sigma^2 \mathrm{i}\phi}, \tag{11}$$

d is as in (6) and $\mathcal{I}_\nu(\cdot)$ is the modified Bessel function of the first kind with ν degrees of freedom.

An application of Lévy's Inversion Theorem similar to the one of Gil-Pelaez [14] leads to the probability distribution function of this random variable.

Corollary 1 (Probability distribution function of the integrated variance over time)**.** *Given $V(t)$ and $V(T)$, the probability distribution function (pdf) of the integrated variance over time under the Heston model is expressed as*

$$f^{\mathbb{Q}}_{IV(t,T)}(z|V(t),V(T)) = \frac{2}{\pi}\int_0^\infty \cos(\phi z)\mathrm{Re}\left(\Xi(\phi)\right)\, d\phi. \tag{12}$$

The formulation provided for the pdf of the integrated variance over time is not computationally friendly. To solve this issue, Glasserman and Kim [15] propose an alternative representation of the integrated variance over time.

Proposition 7 (Gamma representation of the integrated variance over time)**.** *The integrated variance over time in Equation* (9) *admits the representation*

$$IV(t,T) \overset{d}{=} X_1 + X_2 + \sum_{j=1}^{\eta} Z_j$$

where X_1, X_2, $Z_j, \forall\, j$, and η are mutually independent. The random variables X_1, X_2 and Z_j have the following representations:

$$\begin{aligned} X_1 &\overset{d}{=} \sum_{n=1}^{\infty}\frac{1}{\gamma_n}\sum_{j=1}^{N_n} A_j, \\ X_2 &\overset{d}{=} \sum_{n=1}^{\infty}\frac{1}{\gamma_n} B_n, \\ Z_j &\overset{d}{=} \sum_{n=1}^{\infty}\frac{1}{\gamma_n} C_{n,j}, \end{aligned}$$

where

$$\gamma_n = \frac{\kappa^2(T-t)^2 + 4\pi^2 n^2}{2\sigma^2(T-t)^2}. \tag{13}$$

Here, the A_js are independent exponential random variables with mean 1, N_n are independent Poisson random variables with respective means $(V(t)+V(T))\lambda_n$ and

$$\lambda_n = \frac{16\pi^2 n^2}{\sigma^2(T-t)\left[\kappa^2(T-t)^2 + 4\pi^2 n^2\right]}.$$

The B_ns are independent gamma random variables with a shape parameter of $d/2$ and a scale parameter of 1, where d is as in Equation (6)*; $C_{n,j}, \forall n$, are independent gamma random variables with a shape parameter of 2 and a scale parameter of 1. Finally, η is a Bessel random variable with parameter*

$$z = \frac{2\kappa/\sigma^2}{\sinh\left(\kappa(T-t)/2\right)}\sqrt{V(t)V(T)} \tag{14}$$

and $\nu = d/2 - 1$ degrees of freedom.

The moments of this distribution can easily be obtained through the use of a gamma expansion. The expectation and variance of the integrated variance over time shall be employed in the proposed sampling scheme described in Section 4.

Proposition 8 (Moments of the integrated variance over time)**.** *Let* $C_1 = \coth(\kappa\ (T-t)/2)$ *and* $C_2 = \operatorname{csch}^2(\kappa(T-t)/2)$*. Given* $V(t)$ *and* $V(T)$*, the mean and variance of* $IV(t,T)$ *are expressed as*

$$\mathbb{E}[IV(t,T)|V(t),V(T)] = \mathbb{E}[X_1] + \mathbb{E}[X_2] + \mathbb{E}[\eta]E[Z]$$

and

$$\operatorname{Var}[IV(t,T)|V(t),V(T)] = \operatorname{Var}[X_1] + \operatorname{Var}[X_2] + \mathbb{E}[\eta]\operatorname{Var}[Z] + \left(\mathbb{E}[\eta^2] - \mathbb{E}[\eta]^2\right)\mathbb{E}[Z]^2.$$

The mean and variance of X_1 *respectively satisfy*

$$\begin{aligned}\mathbb{E}[X_1] &= [V(t)+V(T)]\left(\frac{C_1}{\kappa} - (T-t)\frac{C_2}{2}\right),\\ \operatorname{Var}[X_1] &= [V(t)+V(T)]\,\sigma^2\left(\frac{C_1}{\kappa^3} + (T-t)\frac{C_2}{2\kappa^2} - (T-t)^2\frac{C_1C_2}{2\kappa}\right);\end{aligned}$$

the mean and variance of X_2 *respectively satisfy*

$$\begin{aligned}\mathbb{E}[X_2] &= d\sigma^2\left(\frac{-2+\kappa(T-t)C_1}{4\kappa^2}\right),\\ \operatorname{Var}[X_2] &= d\sigma^4\left(\frac{-8+2\kappa(T-t)C_1+\kappa^2(T-t)^2C_2}{8\kappa^4}\right);\end{aligned}$$

the mean and variance of Z *respectively satisfy*

$$\mathbb{E}[Z] = 4\mathbb{E}[X_2]/d, \quad \operatorname{Var}[Z] = 4\operatorname{Var}[X_2]/d,$$

with d *as in Equation* (6)*. Finally,*

$$\mathbb{E}[\eta] = \frac{z\mathcal{I}_{d/2}(z)}{2\mathcal{I}_{d/2-1}(z)}, \quad \mathbb{E}[\eta^2] = \frac{z^2\mathcal{I}_{d/2+1}(z)}{4\mathcal{I}_{d/2-1}(z)} + \mathbb{E}[\eta],$$

where z *is as in Equation* (14) *and* $\mathcal{I}_\nu(\cdot)$ *is a Bessel function of the first kind with* ν *degrees of freedom.*

Proof. See [30]. □

3 Simulation Schemes for the Heston Model: A Review

This section aims at describing popular simulation schemes for the Heston model: the generic Euler-Maruyama and Milstein algorithms are first presented, after which Broadie and Kaya's approach [7] is detailed. Several modifications of the latter scheme are available in the literature, and so we describe some of these algorithms: Smith's approximation [28], Broadie and Kaya's drift interpolation of [31], Andersen's quadratic exponential [2], and Tse and Wan's inverse Gaussian [30].

The main goal of these schemes is to simulate asset price paths for a given partition. Let

$$\mathcal{P} = \{0, h, 2h, ..., T\},$$

be an equidistant partition with time step $h \triangleq 1/m$, where m is the number of steps per year. In order to obtain price paths, observations are thus generated for each time value in $\mathcal{P}$.

3.1 Euler-Maruyama scheme

The Euler-Maruyama algorithm, based on a linear approximation of the processes considered, is the most easily implemented method for generating price paths under the Heston model. Generally speaking, it is also one of the most popular schemes as it can be used with virtually any SDE. It consists in a generalization of the Euler method for ordinary differential equations.

Algorithm 3.1 (Euler-Maruyama scheme under the Heston model)**.** *Let $\hat{X}$ and $\hat{V}$ denote discrete-time approximations of X and V respectively. The Euler-Maruyama scheme applied to Equations* (3) *and* (4) *is given by*

$$\hat{X}(hi) = \hat{X}(h(i-1)) + \left[r - \frac{1}{2}\hat{V}(h(i-1))\right]h + \sqrt{\hat{V}(h(i-1)}Z_X\sqrt{h},$$
$$\hat{V}(hi) = \hat{V}(h(i-1)) + \kappa\left[\theta - \hat{V}(h(i-1))\right]h + \sigma\sqrt{\hat{V}(h(i-1))}Z_V\sqrt{h},$$

where Z_X and Z_V are standardized Gaussian random variables such that $\mathrm{corr}(Z_X, Z_V) = \rho$.

Generally speaking, practitioners cannot rely on Algorithm 3.1 to obtain paths of the underlying variance process. When the Feller Condition (Proposition 3) is violated, $\hat{V}(hi)$ may take negative values since the origin is then accessible and strongly reflecting. To avoid negative variances, some modifications are then required.

Lord, Koekkoek and van Dijk [21] propose an important number of solutions to this issue. In particular, their full truncation scheme is a modification of the Euler-Maruyama scheme that is tailored so as to minimize the positive bias on European options.

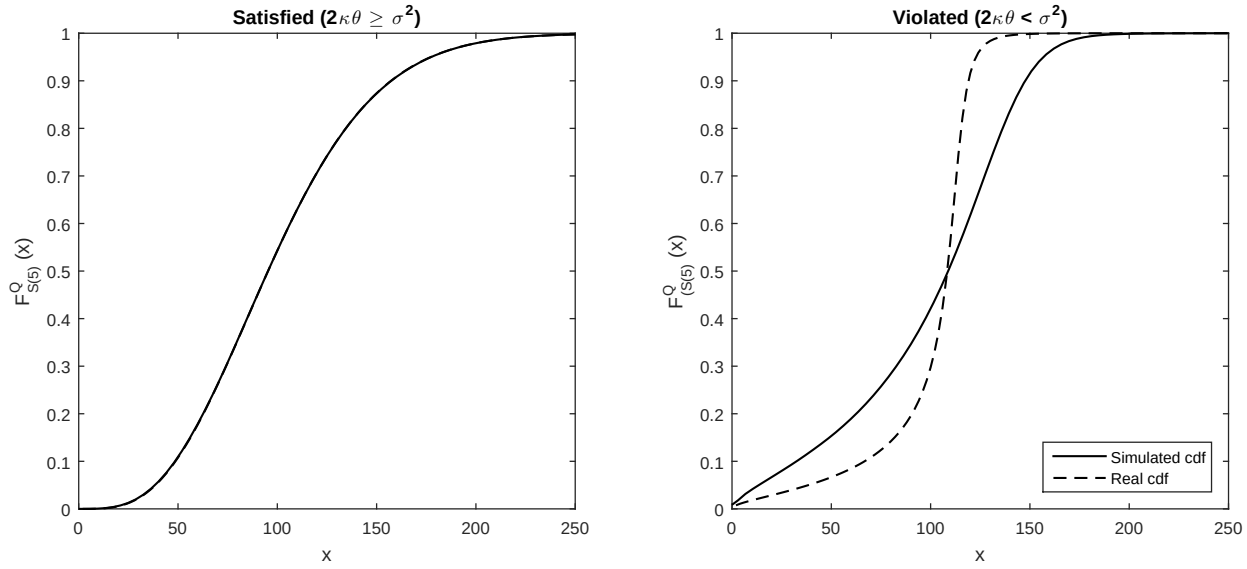

Figure 2: Comparison of a simulated cdf of the price at time 5 using the full truncation scheme (solid line) versus the real cdf of the price at time 5 based on Bakshi et al.'s [4] formulation (dashed line). On the left graph, the Feller Condition is satisfied; on the right graph, it is violated.

Algorithm 3.2 (Euler-Maruyama scheme: Lord et al.'s modification [21])**.** *Under the settings of Algorithm 3.1, let*

$$\hat{X}(hi) = \hat{X}(h(i-1)) + \left[r - \frac{1}{2} f\left(\hat{V}(h(i-1))\right)\right] h + \sqrt{f\left(\hat{V}(h(i-1))\right)} Z_X \sqrt{h}, \tag{15}$$

and

$$\hat{V}(hi) = \hat{V}(h(i-1)) + \kappa\left[\theta - f\left(\hat{V}(h(i-1))\right)\right] h + \sigma \sqrt{f\left(\hat{V}(h(i-1))\right)} Z_V \sqrt{h},$$

where $f(x) = \max(0, x)$.

The full truncation scheme ensures that whenever the $\hat{V}$ process goes below zero it becomes deterministic with an upward drift of $\kappa\theta$. Although this scheme minimizes the positive bias on European call options, it still yields poor results when the Feller Condition is violated; Figure 2 illustrates this issue.

3.2 Milstein scheme

This scheme is similar to the Euler-Maruyama algorithm, but relies on the second-order approximation of a SDE. In practice, the Milstein scheme is applied to the variance process only, as it usually turns out to be the cause of potential problems. In the following algorithm, the full truncation argument is also used in order to avoid negative variances.

Algorithm 3.3 (Milstein scheme with full truncation under the Heston model)**.** *Under the settings*

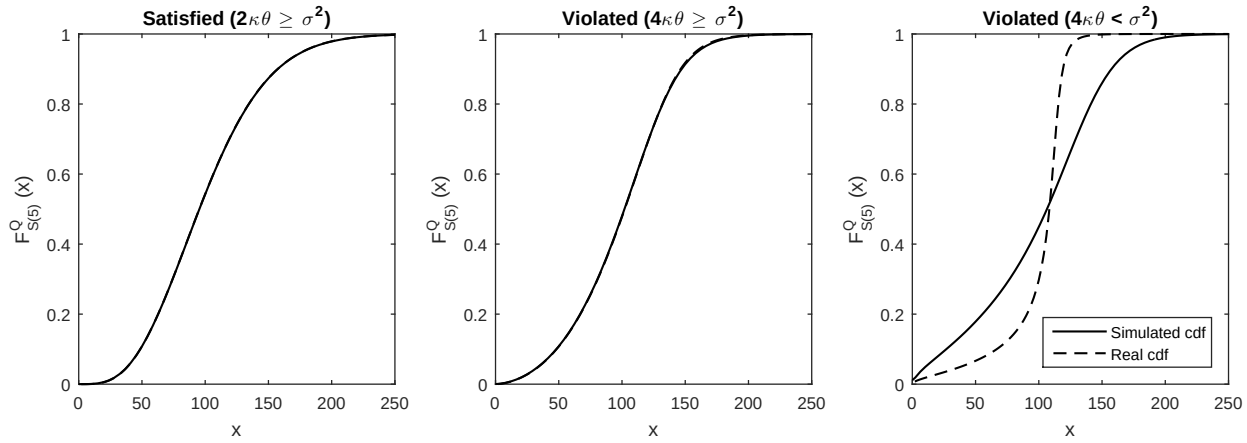

Figure 3: Comparison of a simulated cdf of the price at time 5 using the Milstein scheme with full truncation (solid line) versus the real cdf based on Bakshi et al.'s [4] formulation (dashed line) when the Feller Condition is satisfied (left graph), when the Feller condition is violated but $4\kappa\theta \geq \sigma^2$ (middle graph) and when $4\kappa\theta < \sigma^2$ (right graph).

of Algorithm 3.1, let $\hat{X}$ be as in Equation (15) and $\hat{V}$ satisfy

$$\hat{V}(hi) = \hat{V}(h(i-1)) + \kappa\left[\theta - f\left(\hat{V}(h(i-1))\right)\right]h + \sigma\sqrt{f\left(\hat{V}(h(i-1))\right)}Z_V\sqrt{h} + \frac{\sigma^2}{4}(Z_V^2 - 1)h,$$

where $f(x) = \max(0, x)$.

Gatheral [13] states that if $\hat{V}(h(i-1)) > 0$ and $4\kappa\theta \geq \sigma^2$, then $\hat{V}(hi) > 0$. When this inequality is not satisfied, it still is possible to show that the occurence of negative variances is greatly reduced as compared to the Euler-Maruyama scheme. However, this algorithm still yields poor results when the condition is violated, as illustrated in Figure 3.

3.3 Broadie and Kaya's exact scheme

Broadie and Kaya [7] propose an exact simulation scheme for the Heston model. Even if this method is elegant and theoretically appealing, its practical use is limited. The main issue of this scheme is the lack of speed. A simple simulation using the Euler-Maruyama outperforms this scheme in terms of computational efficiency according to [21].

By applying Itô's Lemma and a Cholesky decomposition to the explicit solution of the asset price in Equation (1),

$$S(hi) = S(h(i-1))\exp\left\{rh - \frac{1}{2}\int_{h(i-1)}^{hi} V(u)\,du + \int_{h(i-1)}^{hi}\sqrt{V(u)}\,dW_S(u)\right\},$$

one obtains an explicit solution for the logarithm of the asset price in (3)

$$X(hi) = X(h(i-1)) + rh - \frac{1}{2}\int_{h(i-1)}^{hi} V(u)\,du + \rho\int_{h(i-1)}^{hi}\sqrt{V(u)}\,dW_V(u) + \sqrt{1-\rho^2}\int_{h(i-1)}^{hi}\sqrt{V(u)}\,dW_X(u), \quad (16)$$

where $W_X(u)$ and $W_V(u)$ are the time-u values of two independent Brownian motions. Integrating the variance process V in Equation (4) yields

$$V(hi) = V(h(i-1)) + \int_{h(i-1)}^{hi}\kappa\left[\theta - V(u)\right]\,du + \sigma\int_{h(i-1)}^{hi}\sqrt{V(u)}\,dW_V(u). \quad (17)$$

By isolating the integral in the last term of the previous equation,

$$\int_{h(i-1)}^{hi}\sqrt{V(u)}\,dW_V(u) = \sigma^{-1}\left[V(hi) - V(h(i-1)) - \kappa\theta h + \kappa\int_{h(i-1)}^{hi}V(u)\,du\right],$$

and then substituting it in (16), one obtains

$$X(hi) = X(h(i-1)) + rh + \frac{\kappa\rho}{\sigma}\int_{h(i-1)}^{hi}V(u)\,du - \frac{1}{2}\int_{h(i-1)}^{hi}V(u)\,du\frac{\rho}{\sigma}\left[V(hi) - V(h(i-1)) - \kappa\theta h\right] + \sqrt{1-\rho^2}\int_{h(i-1)}^{hi}\sqrt{V(u)}\,dW_X(u).$$

The previous equation involves three stochastic quantities that need to be sampled:

1. $V(hi)$ given $V(h(i-1))$;
2. $\int_{h(i-1)}^{hi}V(u)\,du$ given $V(hi)$ and $V(h(i-1))$;
3. $\int_{h(i-1)}^{hi}\sqrt{V(u)}\,dW_V(u)$ given $\int_{h(i-1)}^{hi}V(u)\,du$.

The exact sampling scheme is now described.

Algorithm 3.4 (Broadie and Kaya's exact scheme)**.**

1. *Using Proposition 4, generate $\hat{V}(hi)$ given $\hat{V}(h(i-1))$. To this end, Broadie and Kaya [7] use a result of [18]: see Algorithm 3.5.*
2. *Given $\hat{V}(h(i-1))$ and $\hat{V}(hi)$, generate the integrated variance over time, $\widehat{IV}(h(i-1), hi)$, using the inversion method which requires to invert the cdf*

$$F^{\mathbb{Q}}_{IV(h(i-1),hi)}(x|V(h(i-1)), V(hi)) = \frac{2}{\pi}\int_0^\infty \frac{\sin ux}{u}\mathrm{Re}\left(\Xi(u)\right)\,du$$

where $\Xi(\phi)$ is given in (10)*. This can be done numerically using Newton's second-order method.*

3. *Generate an observation z_X from an independent standardized Gaussian random variable and use the exact solution,*

$$\hat{X}(hi) = \hat{X}(h(i-1)) + rh + \frac{\kappa\rho}{\sigma}\widehat{IV}(h(i-1), hi)$$
$$- \frac{1}{2}\widehat{IV}(h(i-1), hi) + \frac{\rho}{\sigma}\left[\hat{V}(hi) - \hat{V}(h(i-1)) - \kappa\theta h\right] + \sqrt{1-\rho^2}z_X\sqrt{\widehat{IV}(h(i-1), hi)}. \quad (18)$$

To get better computation times, [7] suggest building a cache to store values of IV given several values of $V(hi)$ and $V(h(i-1))$.

In order to generate $\hat{V}(hi)$ given $\hat{V}(h(i-1))$ in Step 1 of Algorithm 3.4, Broadie and Kaya [7] use results from [18] to propose an algorithm for sampling from the non-central chi-squared distribution.

Algorithm 3.5 (Sampling from the non-central chi-squared distribution)**.**

1. *A non-central chi-squared distribution with parameter λ and $\nu > 1$ degrees of freedom satisfies*

$$\chi^2(\nu, \lambda) \overset{d}{=} \left(Z + \sqrt{\lambda}\right)^2 + \chi^2(\nu - 1, 0)$$

where Z is a standardized Gaussian random variable. Thus, if $\nu > 1$, the non-central chi-squared can be generated using $\left(z + \sqrt{\lambda}\right)^2 + c$, where z and c are observations from a standardized Gaussian distribution and a central chi-squared distribution with $\nu - 1$ degrees of freedom, respectively.

2. *When $\nu \leq 1$, a non-central chi-squared random variable may be expressed through a central chi-squared random variable featuring random degrees of freedom. Specifically, if N is a Poisson random variable with mean $\lambda/2$, then $\chi^2(d + 2N, 0) \overset{d}{=} \chi^2(d, \lambda)$. Thus, if $\nu \leq 1$, generate an observation p from a Poisson distribution with mean $\frac{\lambda}{2}$; then, generate an observation from a central chi-squared distribution with $\nu + 2p$ degrees of freedom.*

Algorithm 3.5 requires generating observations from a central chi-squared distribution, which is a special case of the gamma distribution with scale parameter equal to 2. To sample from a gamma distribution, [7] propose using the methods GS* and GKM3 presented in [12]; in this paper, we rely on Marsaglia's method [22]. To generate observations from a Poisson distribution, we use the method expounded in Devroye [9]: if the mean is small, we use the waiting time method; if it is large, we use Ahrens and Dieter's method [1].

The algorithms presented here rely heavily on acceptance-rejection methods. However, we are fully aware that acceptance-rejection methods are ill-suited for quasi Monte Carlo methods and

that these low-discrepancy sequences are widely used in practice.

The acceptance-rejection algorithms presented in this paper are used for simplicity; however, one could employ the direct inversion method for both chi-squared and gamma distribution. This would obviously require precomputation in order to achieve greater efficiency. For instance, van Haastrecht and Pelsser [31] proposed a method to sample from a non-central chi-squared distribution based on direct inversion (so-called NCI). Thus, instead of using Algorithm 3.5, one could use NCI. For the gamma distribution, a similar approach could be implemented, i.e. a cache for quantile function values for different sets of parameters.

While Broadie and Kaya's exact scheme is not directly used in practice, it is the foundation of several sampling algorithms that have been proposed to sample from various stochastic volatility models (and especially the Heston model).

3.4 Smith's approximation

To simplify the implementation of the exact method, Smith [28] uses the similarity between the geometric and arithmetic means of a sample to approximate the characteristic function in Equation (10). By taking a weighted average of these two means,

$$z = \omega \left[\frac{V(h(i-1)) + V(hi)}{2} \right] + (1-\omega) \sqrt{V(h(i-1))V(hi)}$$

where $\omega \in [0,1]$, Smith reduces the number of variables in (10). The resulting approximation of the characteristic function is

$$\begin{aligned}\hat{\Xi}(\phi) = \frac{\gamma(\phi)e^{-\frac{1}{2}(\gamma(\phi)-\kappa)(T-t)}[1-e^{-\kappa(T-t)}]}{\kappa[1-e^{-\gamma(\phi)(T-t)}]} \\ \times \exp\left\{ \frac{2z}{\sigma^2} \left(\frac{\kappa[1+e^{-\kappa(T-t)}]}{1-e^{-\kappa(T-t)}} - \frac{\gamma(\phi)[1+e^{-\gamma(\phi)(T-t)}]}{1-e^{-\gamma(\phi)(T-t)}} \right) \right\} \\ \times \mathcal{I}_{\frac{d}{2}-1}\left(z \frac{4\gamma(\phi)e^{-\frac{1}{2}\gamma(\phi)(T-t)}}{\sigma^2[1-e^{-\gamma(\phi)(T-t)}]} \right) \Big/ \mathcal{I}_{\frac{d}{2}-1}\left(z \frac{4\kappa e^{-\frac{1}{2}\kappa(T-t)}}{\sigma^2[1-e^{-\kappa(T-t)}]} \right),\end{aligned}$$

where $\gamma(\phi)$ and d are as in Equations (11) and (6) respectively. Note that this function depends on $V(h(i-1))$ and $V(hi)$ through z only.

Although faster than the exact scheme of the previous section, this approach still suffers from long computation times as inverting the approximated characteristic function turns out to be a demanding task.

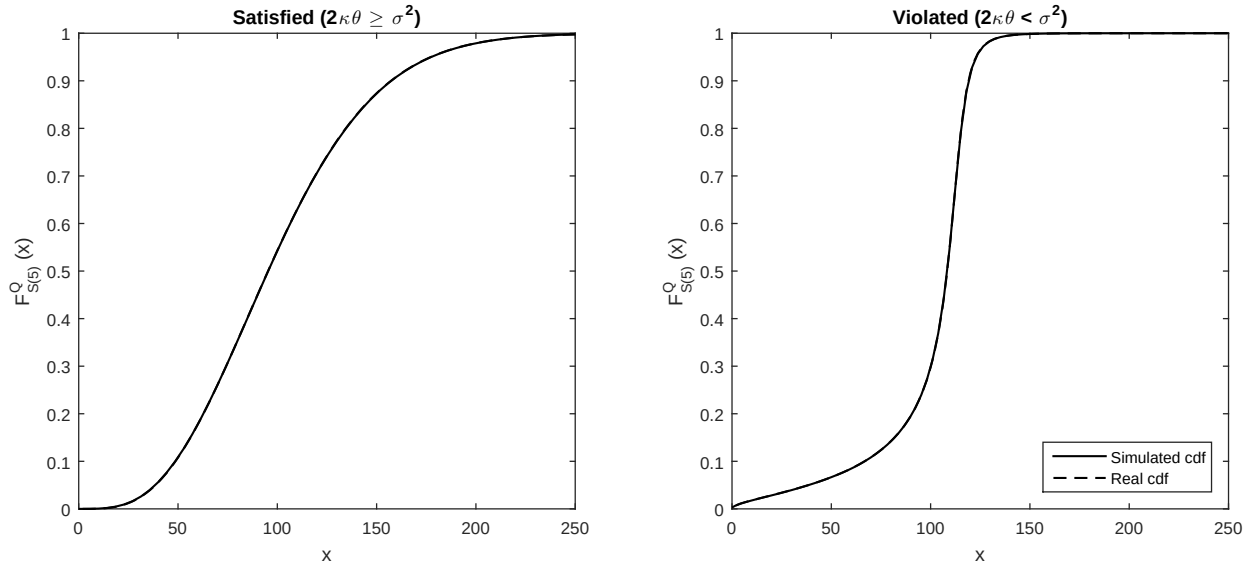

Figure 4: Comparison of a simulated cdf of the price at time 5 using the drift interpolation scheme (solid line) versus the real cdf based on Bakshi et al.'s [4] formulation (dashed line). On the left graph, the Feller Condition is satisfied; on the right graph, it is violated.

3.5 Van Haastrecht and Pelsser's drift interpolation scheme

Van Haastrecht and Pelsser's drift interpolation scheme (BK-I hereafter) [31] aims at speeding up the second step of Algorithm 3.4. To achieve this, $\widehat{IV}(h(i-1), hi)$ is approximated by using a simple quadrature rule called the trapezoidal rule:

$$\widehat{IV}\left(h(i-1), hi\right) \approx h\frac{[V(h(i-1)) + V(hi)]}{2}.$$

The integrated variance over time is thereby deterministic (since there is no sampling involved in this step). Thus, based on this approximation and a small time step, one can obtain good results in seconds. However, this is a coarse approximation when the time step is not extremely small. Figure 4 illustrates two numerical examples of this method when the time step is small.

As stated earlier, van Haastrecht and Pelsser [31] introduce a new method to sample from a non-central chi-squared distribution based on direct inversion. Instead of creating a three-dimensional cache of the inverse of the non-central chi-squared distribution, the authors are able to project it on a one-dimensional search space. This would allow for low-discrepancy numbers.

3.6 Andersen's quadratic exponential scheme

This approach is inspired from the drift interpolation scheme, but instead of using NCI to sample from the non-central chi-squared distribution, Andersen [2] proposes an approximation based on moment-matching techniques. His goal is then to speed up the first step of Broadie and

Kaya [7]'s method.

He observes that a non-central chi-squared random variable with a high non-centrality parameter can be represented by a power function applied to a Gaussian variable. While a cubic transformation of a Gaussian variable is preferable, such a scheme would yield negative values. Accordingly, the quadratic representation of Patnaik [27] is preferred. For a sufficiently large value of $\hat{V}(h(i-1))$, we get

$$\hat{V}(hi) = a(b + Z_V)^2, \tag{19}$$

where Z_V is a standardized Gaussian random variable, and a and b are functions of $\hat{V}(h(i-1))$, whose values are determined by moment-matching techniques.

According to [20], $\hat{V}(hi)$ is distributed as a times a non-central chi-squared distribution with one degree of freedom and non-centality parameter b^2. Then,

$$\mathbb{E}\left[\hat{V}(hi)\right] = a(1 + b^2) \quad \text{and} \quad \text{Var}\left[\hat{V}(hi)\right] = 2a^2(1 + 2b^2). \tag{20}$$

Equating each moment in (20) to the exact values

$$m \triangleq \mathbb{E}\left[\hat{V}(hi)|\hat{V}(h(i-1))\right] \quad \text{and} \quad s^2 \triangleq \text{Var}\left[\hat{V}^2(hi))|\hat{V}(h(i-1))\right]$$

respectively (see Subsection 2.3) yields

$$a = \frac{m}{1 + b^2} \quad \text{and} \quad b^2 = 2\psi^{-1} - 1 + \sqrt{2\psi^{-1}}\sqrt{2\psi^{-1} - 1}, \tag{21}$$

where $\psi = \frac{s^2}{m^2}$ if $\psi \leq 2$.

Equation (19) does not work well for small values of $\hat{V}(h(i-1))$. In such cases, to get a better approximated density for $\hat{V}(hi)$, Andersen [2] proposes

$$\mathbb{Q}(\hat{V}(hi) \in [x, x + dx]|\hat{V}(h(i-1))) \approx \left[p\delta_0(x) + \beta(1-p)e^{-\beta x}\right]dx, \tag{22}$$

where $x \geq 0$, $\delta_0(\cdot)$ is a Dirac delta function at the origin, and p and β are non-negative constants to be determined. The resulting density has a probability mass at the origin, supplemented with an exponential tail. Sampling from Equation (22) using the inverse method is straightforward since the associated cdf is easily invertible:

$$\Psi(x; p, \beta) \triangleq \hat{F}^{\hat{\mathbb{Q}}}_{\hat{V}(hi)}(x|\hat{V}(h(i-1))) = p + (1-p)(1 - e^{-\beta x}), \quad x \geq 0. \tag{23}$$

The constants p and β are derived using, once again, a moment-matching argument. Integrating (22) leads to

$$\hat{\mathbb{E}}\left[\hat{V}(hi)\right] = \frac{1-p}{\beta} \quad \text{and} \quad \widehat{\text{Var}}\left[\hat{V}(hi)\right] = \frac{1-p^2}{\beta^2}.$$

Using the theoretical moments, one obtains for $\psi \geq 1$,

$$p = \frac{\psi - 1}{\psi + 1} \quad \text{and} \quad \beta = \frac{1 - p}{m}. \tag{24}$$

Using the appropriate part of the scheme requires the creation of a switching rule. Equations (19) and (21) may be used when $\psi \leq 2$, while (23) and (24) are used when $\psi \geq 1$. These domains of applicability overlap which means that at least one of these schemes can be used; one may thus introduce a critical level $\psi_c \in [1, 2]$ as a switching rule. Any ψ_c in this interval is appropriate with the QE scheme.

For the process X, Andersen uses the following discretization

$$\hat{X}(hi) = \hat{X}(h(i-1)) + rh + K_0 + K_1\hat{V}(h(i-1)) + K_2\hat{V}(hi) + \sqrt{K_3\left[\hat{V}(h(i-1)) + \hat{V}(hi)\right]}Z_X, \tag{25}$$

with

$$K_0 = -\frac{\rho\kappa\theta}{\sigma}h, \qquad K_1 = \frac{h}{2}\left(\frac{\kappa\rho}{\sigma} - \frac{1}{2}\right) - \frac{\rho}{\sigma},$$
$$K_2 = \frac{h}{2}\left(\frac{\kappa\rho}{\sigma} - \frac{1}{2}\right) + \frac{\rho}{\sigma}, \qquad K_3 = \frac{h}{2}(1 - \rho^2).$$

We outline that fact that the constants K_0, K_1, K_2 and K_3 described above have been simplified to be similar to what is already presented in Equation (18); for general formulations of these constants, see [2].

Algorithm 3.6 (Quadratic exponential scheme)**.**

1. *Compute m and s^2 from $\hat{V}(h(i-1))$.*
2. *If $\psi = \frac{s^2}{m^2} \leq \psi_c$,*
   (a) *compute a and b;*
   (b) *draw z_V, an observation from a standardized Gaussian random variable;*
   (c) *compute $\hat{V}(hi) = a(b + z_V)^2$.*
3. *If $\psi > \psi_c$,*
   (a) *compute β and p;*
   (b) *draw u_V, an observation from a uniform random variable on $(0, 1)$;*
   (c) *compute $\hat{V}(hi) = \Psi^{-1}(u_V; p, \beta)$.*
4. *Draw z_X, an observation from an independent standardized Gaussian random variable.*

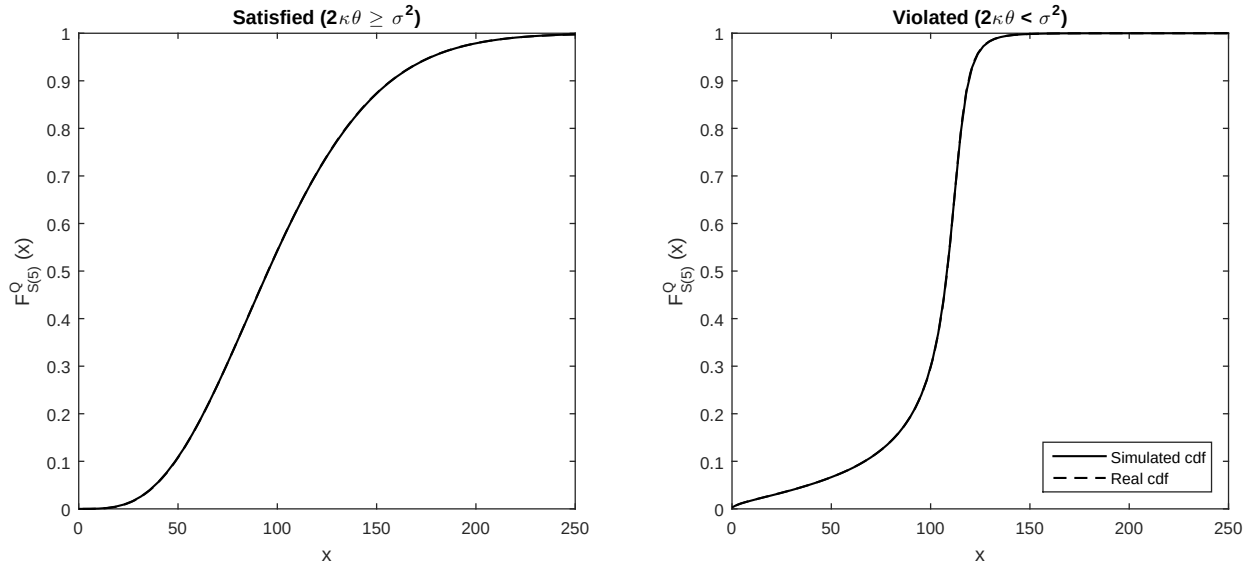

Figure 5: Comparison of a simulated cdf of the price at time 5 using the quadratic exponential scheme (solid line) versus the real cdf based on Bakshi et al.'s [4] formulation (dashed line). On the left graph, the Feller Condition is satisfied; on the right graph, it is violated.

5. *Compute* $\hat{X}(hi)$ *based on* $\hat{V}(h(i-1))$, $\hat{V}(hi)$, $\hat{X}(h(i-1))$ *and* z_X, *as in BK-I.*

This sampling scheme generally yields good results, even when the Feller Condition is violated (see Figure 5).

3.7 Tse and Wan's inverse Gaussian scheme

Tse and Wan's inverse Gaussian (hereafter, IG) scheme [30], which relies on the basics of Broadie and Kaya [7], uses the fact that the integrated variance over time converges in distribution to a moment-matched inverse Gaussian random variable as h goes to infinity.

Simulating observations from an inverse Gaussian random variable is accessible; we refer the reader to [24] for more detail. To increase the speed of their algorithm, Tse and Wan however propose a modification to Michael, Schucany and Haas' algorithm.

Algorithm 3.7 (Sampling from an inverse Gaussian distribution)**.** *To generate observations from an inverse Gaussian distribution with mean* μ *and shape parameter* λ*:*

1. *Generate an observation* z *from a standardized Gaussian distribution.*

2. *Let*
$$x = 1 + \frac{z^2}{2\lambda} - \frac{1}{2\lambda}\sqrt{4\lambda z^2 + z^4}.$$

3. *Generate an observation* u *from a uniform distribution on* $(0, 1)$.

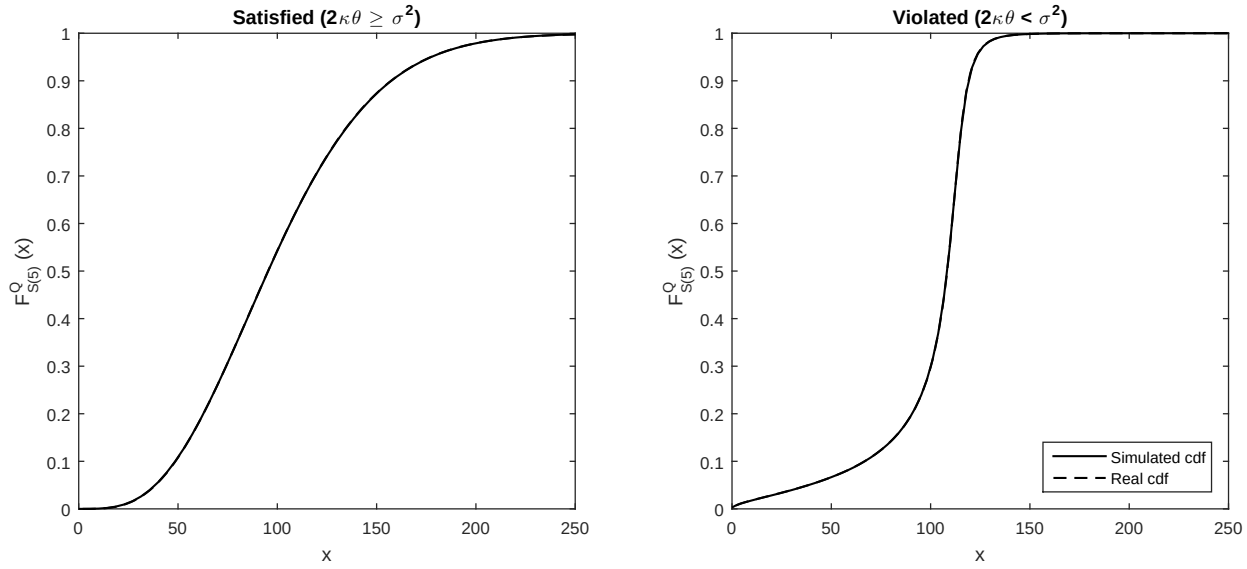

Figure 6: Comparison of a simulated cdf of the price at time 5 using the inverse Gaussian scheme (solid line) versus the real cdf based on Bakshi et al.'s [4] formulation (dashed line). On the left graph, the Feller Condition is satisfied; on the right graph, it is violated.

4. If $u \leq 1/(1+x)$, then return μx; otherwise, return μ/x.

Proof. See [24] and [30]. □

Tse and Wan [30] also propose a new algorithm to generate a value of $V(hi)$ given $V(h(i-1))$. The IPZ scheme is based on precomputation of quantile function values for some cases that could take more time to be generated. For the other cases, the authors used something similar to Algorithm 3.5.

Figure 6 illustrates the results of a simulation study using the IG scheme (without implementing the IPZ scheme). This approach seems to yield good results, even when the Feller Condition is violated.

4 Gamma Approximation Scheme

This section discloses the details about the gamma approximation (hereafter GA) algorithm for simulating from the Heston model. The basics of the scheme are first discussed, after which a formal version of the algorithm is presented. The caches used in this scheme are finally described.

4.1 Approximating the integrated variance over time

As several of the algorithms discussed in Section 3, the GA algorithm focuses on generating an observed integrated variance over time so as to improve the computational efficiency of the exact algorithm while eliminating the issues inherent to the variants previously discussed. The resulting approach is thus a modification of Broadie and Kaya's exact algorithm. Note that the first step of Algorithm 3.4 could also be improved, but this matter is not pursued here.

As detailed in Proposition 7, Glasserman and Kim [15] propose an alternative representation for the integrated variance over time defined in (9). Furthermore, Tse and Wan [30] use Laplace transforms to provide the first two moments of this random variable; see Proposition 8. In their paper, they also demonstrate the convergence in distribution of the integrated variance over time towards a moment-matched inverse Gaussian as $h \to \infty$, and use this argument to justify the use of this distribution in their proposed algorithm (see Subsection 3.7). The validity of this justification is however questionable, as small time steps h are more appealing for pricing some path-dependent securities (i.e. Bermudian options, lookback options, among others). Accordingly, a convergence result available for large h does not necessarily seem appropriate in such cases.

One may find it more convincing to study the distribution of the integrated variance over time for small values of h. In fact for small h, it is easily verified that γ_n in (13),

$$\gamma_n = \frac{\kappa^2 h^2 + 4\pi^2 n^2}{2\sigma^2 h^2}$$

diverges rapidly as n increases. Indeed, $\gamma_n < \gamma_{n+1} \Leftrightarrow n^2 < (n+1)^2$ and $1/\gamma_n$ tends to zero rapidly. Consequently, the mixtures X_1, X_2 and Z_j of Proposition 7 can be accurately approximated by a finite sum of independent gamma random variables. The exact expression for the convolution of independent gamma distributions with different scale parameters is nontrivial. According to Stewart et al. [29], one may approximate this complex distribution by a simple moment-matched gamma distribution. Visual inspection (see Figure 7) corroborates this affirmation, and thus a gamma distribution seems like a good approximation for the integrated variance over time. The inverse Gaussian betrays a lack of precision when h tends to zero.

Practically speaking, the use of a cache shall allow to efficiently determine the moment-matched parameters of the gamma distribution at every iteration.

4.2 Path simulation of the asset price process

We now introduce the gamma approximation scheme.

Algorithm 4.1 (Gamma approximation scheme)**.**

1. *Create caches for the moments of* $IV(h(i-1), hi)$*; Subsection 4.3 details the caching method*

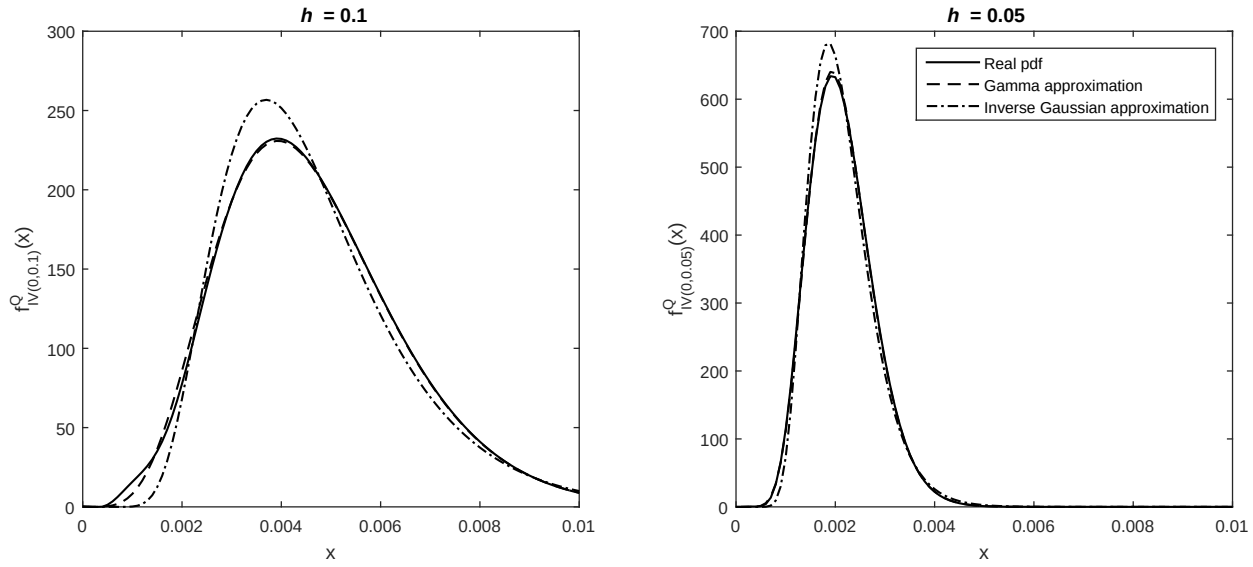

Figure 7: Comparison between the inverse Gaussian approximation (dashed-dot line), the gamma approximation (dashed line), and the real pdf (solid line) of the integrated variance over time, for a small time step; $h = 0.1$ in the left graph, $h = 0.05$ in the right graph.

implemented for this scheme.

2. *Using Proposition 4 and Algorithm 3.5, generate* $\hat{V}(hi)$ *given* $\hat{V}(h(i-1))$.

3. *Given* $\hat{V}(hi)$ *and* $\hat{V}(h(i-1))$*, generate the integrated variance over time,* $\widehat{IV}(h(i-1), hi)$*, from a moment-matched gamma distribution using the moments available in the caches.*

4. *Generate* z_X*, an observation from an independent standardized Gaussian random variable, and use the exact solution in* (18) *to obtain* $\hat{X}(hi)$.

4.3 Caches implementation

We now describe a method to obtain the moments of $IV(h(i-1), hi)$. This task is crucial as the approximated gamma distribution used to generate observations from the integrated variance over time depends on these moments.

Definition 1. *Let* $IV^*(h(i-1), hi) \triangleq IV(h(i-1), hi) - X_1$*, where* $IV(h(i-1), hi)$ *and* X_1 *are given in Proposition 7. From Proposition 8,*

$$\mathbb{E}[IV^*(h(i-1), hi) | V(h(i-1)), V(hi)] = \mathbb{E}[X_2] + \mathbb{E}[\eta] E[Z]$$

and

$$\text{Var}\left[IV^*(h(i-1), hi) | V(h(i-1)), V(hi)\right] = \text{Var}[X_2] + \mathbb{E}[\eta]\ \text{Var}[Z] + \left(\mathbb{E}[\eta^2] - \mathbb{E}[\eta]^2\right)\mathbb{E}[Z]^2.$$

The previous two moments depend on $V(h(i-1))$ and $V(hi)$ through the product $V(h(i-1)) \times V(hi)$ only. Again from Proposition 8,

$$\mathbb{E}[IV(h(i-1), hi)|V(h(i-1)), V(hi)] = \mathbb{E}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)] + \mathbb{E}[X_1]$$

and

$$\text{Var}[IV(h(i-1), hi)|V(h(i-1)), V(hi)] = \text{Var}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)] + \text{Var}[X_1].$$

Since $\mathbb{E}[X_1]$ and $\text{Var}[X_1]$ depend only on $V(h(i-1)) + V(hi)$ and do not require any modified Bessel evaluation, it is thus computationally inexpensive to compute these quantities at a later moment.

Proposition 9 (Moments computations)**.** *For fast moments computations using a cache:*

1. *Precompute* $\mathbb{E}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)]$ *and* $\text{Var}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)]$ *using a* ***special grid*** *(see Definition 2) containing specific values for* $V(h(i-1)) \times V(hi)$*. These values are called the caches.*

2. *Compute* $\mathbb{E}[X_1]$ *and* $\text{Var}[X_1]$.

3. *Use linear interpolation to approximate* $\mathbb{E}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)]$ *and* $\text{Var}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)]$ *from their respective cache.*

4. *Add* $\mathbb{E}[X_1]$ *and* $\text{Var}[X_1]$ *to the previous moments respectively to obtain* $\mathbb{E}[IV(h(i-1), hi)|V(h(i-1)), V(hi)]$ *and* $\text{Var}[IV(h(i-1), hi)|V(h(i-1)), V(hi)]$.

Step 1 of the previous proposition refers to the need of building the caches, i.e. determining values $V(h(i-1)) \times V(hi)$ at which the moments of $IV^*(h(i-1), hi)$ should be evaluated. It turns out that when $\mathbb{E}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)]$ and $\text{Var}[IV^*(h(i-1), hi)|V(h(i-1)), V(hi)]$ are regarded as functions of $V(h(i-1)) \times V(hi)$, their graphs look like piecewise linear curves on a log-log scale. Figure 8 illustrates this behaviour. This suggests using a grid as described below.

Definition 2. *To build the moments' caches, a naturally arising grid for the values* $V(h(i-1)) \times V(hi)$ *is*

$$\left\{0, \delta^{-\zeta}, \delta^{-\zeta+1}, \delta^{-\zeta+2}, \ldots, \delta^{\xi-1}, \delta^{\xi}\right\} \tag{26}$$

where δ *is a constant selected in* $(1, \infty)$ *and* ζ, ξ *are constants in* $\mathbb{N}$.

Using such an exponentially spaced grid allows caching a representative behaviour of the moments of $IV^*(h(i-1), hi)$. As mentioned earlier, the expectation and variance exhibit an almost linear behaviour on a log-log scale, meaning that in reality, these moments display an exponential behaviour.

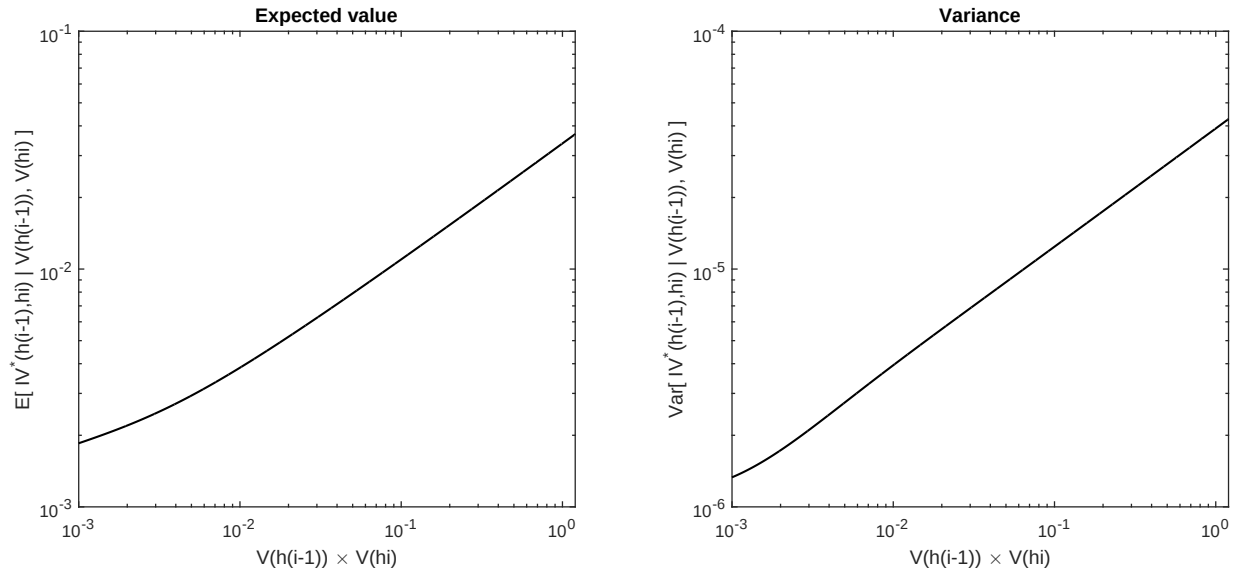

Figure 8: Log-log graph of $\mathbb{E}[IV^*(h(i-1),hi)|V(h(i-1)),V(hi)]$ (left) and $\text{Var}[IV^*(h(i-1),hi)|V(h(i-1)),V(hi)]$ (right) as a function of $V(h(i-1)) \times V(hi)$.

This grid is somewhat different than the one used by Tse and Wan [30]: the authors used an equidistant spaced grid for $\sqrt{V(h(i-1)) \times V(hi)}$. Moreover, they grid seems to require a lot of points (i.e. $2^{15+ceil(log_2(N))}+1$). Ours requires far less points and is very efficient.

It could be argued that a clever type of interpolation, rather than a linear interpolation, would be more appropriate to approximate the moments at a value $V(h(i-1)) \times V(hi)$ not explicited in the cache. In addition to be unecessary, this approach would be significantly more intensive computationally. A simple linear interpolation computationally outperforms the other types of interpolation in this case (the exactness acquired is not significant and the time needed to perform such an algorithm is momentous). Cubic spline and shape-preserving piecewise cubic interpolation were used for comparison.

5 Numerical Study

5.1 Description of the scenarios

In order to evaluate the gamma approximation scheme, we define six cases that are described in Table 1. The first three cases correspond to those introduced in Andersen [2]. In all these scenarios, the Feller Condition (see Proposition 3) is violated. The main difference between these cases is the strike price. The first one corresponds to an "at-the-money" option ($K = 100$), the second one to an "out-of-the-money" option ($K = 140$) and the last one to an "in-the-money" option ($K = 70$). The following two cases (4 and 5) describe a volatile asset featuring a strong underlying correlation structure. The first one has a maturity of ten years and the second a maturity

of five years. The last scenario (Case 6) is studied in Smith [28]; in this case, the Feller Condition is satisfied.

We price two types of options. The first type is a vanilla European call option that pays $\max(S(T) - K, 0)$ at maturity; the second type is an Asian call payout option, which is defined below. The "exact" prices of the European and Asian options for the various scenarios considered may be found in Table 1.

Definition 3 (Asian call payout option)**.** *Let $S(t)$ be the price of an asset at time t. The payout at time T of the Asian option considered in this paper is*

$$A\left(\{S(t)\}_{t\geq 0}, K, T\right) = \max\left(M(T) - K, 0\right)$$

where K is the strike price and $M(T)$ is based on the arithmetic mean, i.e.

$$M(T) = \frac{1}{T}\sum_{i=1}^{T} S(i)$$

here, $S(i)$ is the price of the asset at the end of the i-th year. Thus, the price of this option at time 0 is

$$\mathbb{E}\left[e^{-rT} A\left(\{S(t)\}_{t\geq 0}, K, T\right)\right] \tag{27}$$

The computation times are obtained by running MATLAB programs on a desktop computer with an 4.5 GHz Intel Core i7-2600K and 16 GB RAM. Three sampling schemes are used for the comparison: Andersen's quadratic exponential (QE), van Haastrecht and Pelsser's drift interpolation (BK-I) and Tse and Wan's inverse Gaussian (IG). The Euler-Marayuma and Milstein schemes are discarded since they do not yield good results when some conditions are violated.

The exact prices of European call options are obtained using Bakshi et al.'s [4] formulation. The "exact" prices of Asian options are obtained with the GA scheme using $N = 2^{30}$ paths and $m = 32$ steps per year. We also used the other schemes to see whether this price is robust. There were virtually no difference between the prices given by each other method.

The parameters used to compute the moments' caches are $\zeta = 52$,

$$\xi = \frac{\log(10\sigma)}{\log(\delta)},$$

and

$$\delta = \exp\left(\frac{\log(10\sigma)}{(75-\zeta) + \log_2(N)}\right).$$

They yield accurate results in a reasonable time and they are dependent on the on the volatility of V parameters. Moreover, δ and ξ depend on N to keep the precomputation time proportional to the total simulation time as N varies. The new grid seems to yield good results for a large selection

of these parameters. The efficiency of the method does not seem to be affected by a change in δ, ζ and ξ; though, they need to be large enough.

For each method and each case, the price is obtained for different numbers of steps per year. We modify the number of steps per year, m, by keeping the number of paths constant at 2^{23}.

Table 1: Heston model parameters for Cases 1 to 6.

Case	1	2	3	4	5	6
σ	1	1	1	1	1	0.5196
κ	0.5	0.5	0.5	1	1	1.0407
θ	0.04	0.04	0.04	0.04	0.04	0.0586
$V(0)$	0.04	0.04	0.04	0.04	0.04	0.0194
ρ	-0.9	-0.9	-0.9	-0.999	-0.999	-0.6747
r	0	0	0	0	0	0
T	10	10	10	10	5	4
$S(0)$	100	100	100	100	100	100
K	100	140	70	100	100	100
Exact price (European)	13.0847	0.2958	35.8498	9.5868	6.2001	5.5871
"Exact" price (Asian)	8.1941	0.0243	32.6236	9.9975	7.2683	9.7103

[1] The exact prices of European call options are obtained using Bakshi et al.'s [4] formulation.
[2] The "exact" prices of Asian call options are obtained with the GA scheme using $N = 2^{30}$ paths and $m = 32$ steps per year. We also used the other schemes to see whether this price is robust.

5.2 European option

Let α be the exact price of an option and α' the estimator returned by a simulation scheme. Then, the bias of the estimator (or, more precisely, the numerical error) is given by $(\mathbb{E}[\alpha'] - \alpha)$ and the standard deviation is given by $\sqrt{\mathbb{E}[(\alpha' - \mathbb{E}[\alpha'])^2]}$.

In order to assess our new GA scheme and compare it with the three other schemes (QE, BK-I and IG), we estimate these performance metrics by using Monte Carlo simulations to estimate $\mathbb{E}[\alpha']$ and by calculating the sample standard deviation.

The standard deviation divided by the square root of the sample size (in our case $N = 2^{23}$) helps us measuring how statistically significant the numerical error estimate is. In the figures, we draw horizontal lines to show the levels of one, three and five standard deviations divided by the square root of the sample size. Note that our standard deviations are computed with the GA scheme when $m = 1$. Within each parametric case, the standard deviations do not change substantially between the schemes and the choice of m.

Panel A of Tables 2 provides relative biases and computation times for the first case. Results for Cases 2 to 6 are given in Tables 3 and 4 of Appendix A. The relative biases for which their absolute counterparts are not statistically different from zero are in bold in the tables. In other words, absolute biases that are less than three standard deviations divided by the square root of N are considered statistically nonsignificant. The number of steps were chosen in order to obtain similar computational times among the schemes. However, sometimes, it was difficult to yield

comparable computational times; in those cases, we use larger m values for QE, BK-I and IG. For example, we use a partition going from 0.4 to 2.0 (by steps of 0.1) for GA in the first case. However, we use 0.6 to 2.2 (by steps of 0.1) for IG, 0.8 to 4.0 (steps of 0.2) for BK-I and 1.6 to 8.0 (steps of 0.4) for QE. The first, fourth, seventh and tenth columns of Tables 2, 3, and 4 contain the different m values for each case and each scheme.

The time needed to pre-compute the caches were included in all the timing presented throughout this paper. However, one could store the caches to yield even better results for IG and GA (if he is using the same model's parameters).

In most cases, the numerical errors yielded by these other approaches were greater than the one obtained by GA. Thus, for these occurrences, GA appears to be more efficient than the other schemes in terms of numerical errors and computational times.

Essentially, the proposed scheme performs well when pricing European call options. The results from our GA scheme are similar or better than what is obtained using QE, BK-I and IG schemes according to Tables 2, 3 and 4. In many cases, the GA scheme performs better than the other schemes. For the second case where T is large and K is 140 ("out-of-the-money" option), it is not clear which scheme yields the best results. IG and GA seem to be quite similar in this case.

Figure 9 illustrates the relationships between absolute numerical errors and computational times for the four sampling schemes mentioned above on a log-log scale. Absolute numerical errors are now used to be able to compare them efficiently with standard deviations computed earlier. Each dot in a given curve corresponds to a simulated price bias. By varying m, absolute price biases are obtained for a given scheme (each corresponding to a dot in a given curve). The number of steps per year m varies across the different schemes to take account of the fact that BK-I and QE are faster than the other two methods (but not necessarily better). The simulated price biases in a given curve were all obtained according to a common number of simulated paths N. When we compare the four schemes with larger values of m, the comparison is more tricky. Graphically, the curves oscillate as the number of steps per year becomes larger (mainly for GA and IG). Fortunately, these cases represent bias levels so low that they are likely to be accurate enough in practice.

All six cases represented in Figure 9 show very convincing results. In general, even if the number of steps is not very large, GA yields small biases on European options and leads to decreasing biases as m increases.

Normally, if we witness a bias less than 3 times the standard deviation, it is assumed to be statistically nonsignificant (just like we did previously). Thus, our scheme produces numerical errors that are statistically nonsignificant for almost all scenarios. Case 2 seems to be problematic in the way that it is the only one where IG performs slightly better than GA. However, the price of this option is only 30 cents, which means that an error of 0.001 for the GA scheme is not serious

Table 2: Results for the European call option and Asian call option using Case 1.

Panel A: Case 1, European call option

QE			GA			IG			BK-I		
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
1.6	3.836	18.79	0.4	2.660	17.94	0.6	0.762	21.99	0.8	2.551	18.54
2	2.511	23.07	0.5	1.464	22.34	0.7	0.674	25.86	1	1.984	23.72
2.4	1.723	28.22	0.6	0.757	26.78	0.8	0.472	29.39	1.2	1.525	28.15
2.8	1.189	33.00	0.7	0.425	31.15	0.9	0.363	32.98	1.4	1.235	33.02
3.2	0.878	37.59	0.8	0.268	35.61	1	0.347	36.82	1.6	1.071	37.92
3.6	0.661	42.43	0.9	0.182	40.07	1.1	0.284	40.50	1.8	0.841	42.95
4	0.399	47.09	1	**0.063**	44.47	1.2	0.247	44.23	2	0.665	47.92
4.4	0.410	51.10	1.1	**0.075**	48.97	1.3	0.140	47.99	2.2	0.531	52.92
4.8	0.256	55.80	1.2	**0.094**	53.64	1.4	**0.099**	51.81	2.4	0.492	58.01
5.2	0.193	60.45	1.3	**0.005**	57.95	1.5	**0.067**	55.55	2.6	0.453	63.14
5.6	0.120	65.69	1.4	**0.091**	64.72	1.6	**0.080**	60.22	2.8	0.363	67.94
6	**0.055**	69.41	1.5	**0.002**	66.63	1.7	0.161	63.06	3	0.327	73.85
6.4	**0.007**	75.05	1.6	0.106	72.44	1.8	**0.046**	67.21	3.2	0.300	78.87
6.8	**0.034**	79.43	1.7	**0.029**	75.97	1.9	**0.004**	70.92	3.4	0.282	84.14
7.2	**0.026**	85.62	1.8	**0.068**	81.07	2	**0.082**	74.99	3.6	0.256	88.47
7.6	**0.052**	87.81	1.9	**0.032**	84.84	2.1	**0.019**	78.26	3.8	0.189	94.17
8	**0.025**	92.62	2	**0.028**	89.19	2.2	**0.019**	82.21	4	0.172	99.49

Panel B: Case 1, Asian call option

QE			GA			IG			BK-I		
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
4	0.284	45.97	1	0.216	44.59	1	0.413	36.50	2	0.220	47.85
8	0.118	95.50	2	**0.013**	89.65	2	**0.067**	75.49	4	**0.073**	102.06
12	**0.100**	272.82	3	**0.007**	159.35	3	**0.006**	115.40	6	**0.009**	329.79
16	**0.059**	344.24	4	**0.006**	219.24	4	**0.010**	153.18	8	**0.010**	518.82

[1] The relative biases are in percent.
[2] Relative biases in bold show that their absolute counterparts are statisically nonsignificant (i.e. absolute bias less than three time the standard deviation).

compared to the actual price of the option!

When the relative biases of IG are smaller than the ones produced by GA, the errors resulting from the gamma approximation scheme are quite close to the ones produced by IG. Thus, again, our scheme yields more than adequate results.

When m is small, IG seems to slightly outperform our new scheme. This is not surprising, as the IG scheme has been designed so that the approximated integrated variance over time's behaviour be close to the real one when the time step is large (Tse and Wan's main result).

5.3 Asian option

Panel B of Table 2 provides relative numerical errors and computational times for Case 1. Results for Cases 2 to 6 are given in Table 5 of Appendix A. Values in bold still mean that the numerical errors are less than three standard deviations divided by the square root of the number of paths N.

Again, the numbers of steps were chosen in order to obtain similar computational times among the schemes. Nonetheless, it was even more difficult in the Asian option case than it was with the European option case. Our numbers of steps per year needed to be integers since the payoff of the

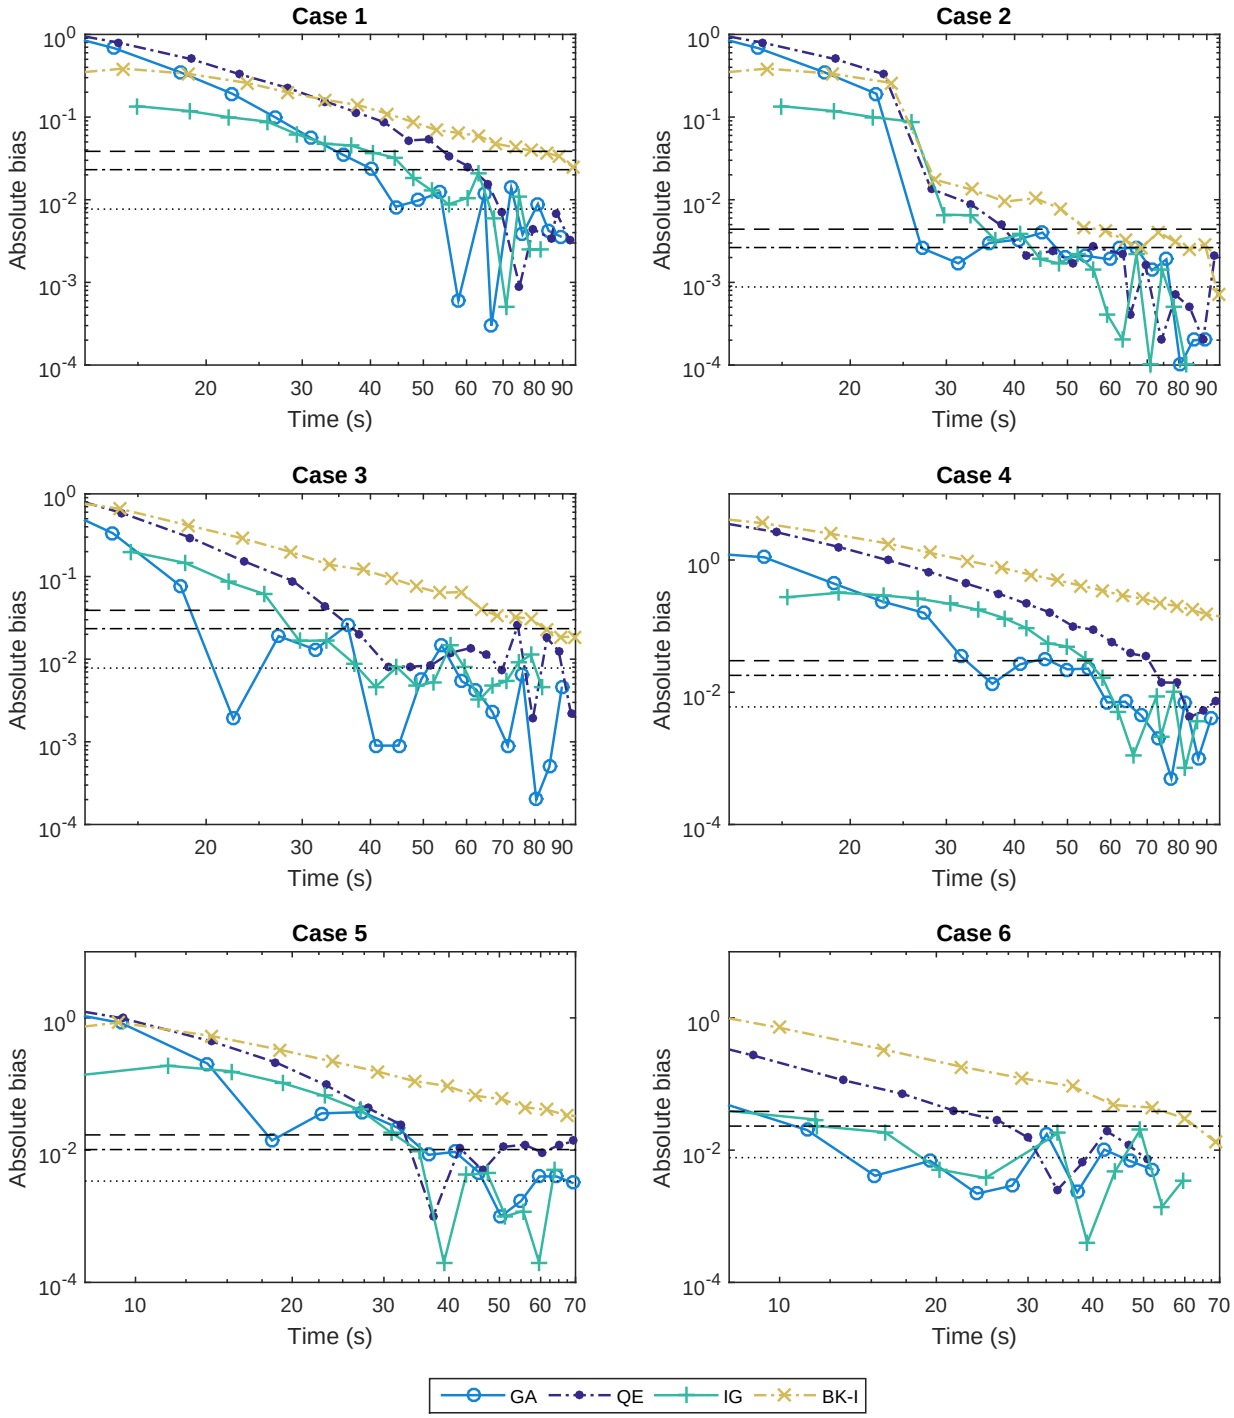

Figure 9: Absolute bias against computational time on pricing European options. The light blue line (circle marker) is used for the GA scheme, navy blue (dot marker) for QE, yellow (cross marker) for BK-I, and turquoise (plus marker) for IG. Standard deviation of the price (dotted line), 3 times the standard deviation (dashed-dot line), and 5 times the stardard deviation (large-dash line) are included.

option depends on the arithmetic mean of the prices of the asset at the end of each year.

Essentially, our scheme seems to perform efficiently when pricing Asian call options. For all scenarios, the results from our GA scheme are better (or, at worst, equal) to QE, BK-I, and IG methods. Note that Case 2 does not appear to be problematic anymore for the Asian option. The numerical errors generated by GA are smaller than one standard deviation divided by the square root of N when $m = 3$ and $m = 4$. When $m = 2$, it is not statistically different from zero for Cases 1, 2, 3, 4, and 6.

Figure 10 is the equivalent of Figure 9 for Asian options; however, we have considered only one data point for GA ($m = 3$) and IG ($m = 4$) in order to show how little is the bias for these two when compared to QE and BK-I. All the scenarios considered in these graphs depict convincing results as far as the GA method is concerned.

According to our tables, when m becomes larger, the results produced by GA seem to be consistent with those obtained with QE, BK-I, and IG. The biases are generally lower than one standard error, thus almost every bias is statistically nonsignificant.

Again, when the relative numerical errors of QE, BK-I, and IG are smaller than the ones produced by GA, the errors resulting from the gamma approximation scheme are quite close to the ones produced by other methods.

6 Concluding Remarks

This paper introduces a new method to simulate stock price paths under the Heston framework. When the Feller Condition is violated, generic schemes generally fail to provide accurate results. After an exhasutive review of available simulation schemes under the Heston model, the gamma approximation algorithm is introduced. This method is based on Stewart et al. [29]: using a moment-matched gamma distribution to sample observations for the integrated variance over time, price paths can be found in an efficient way. This new algorithm is quick and often appears to be slightly more efficient than the most popular schemes for sampling under the Heston model.

Simulation studies demonstrate that the GA scheme performs as well as the most popular schemes available in the literature. In several cases, the GA algorithm seems to outperform QE, BK-I, and IG (for the same computational times). In the remaining cases, the relative biases obtained are close to those witnessed with QE, BK-I, and IG. When m becomes larger, the numerical errors tend to become smaller for the GA scheme than for the other schemes considered.

A deeper analysis of the integrated variance over time distribution could lead to another approximation; the simple gamma approximation however appears to be a good compromise between accuracy and efficiency. The use of gamma mixtures or convolutions could provide a better

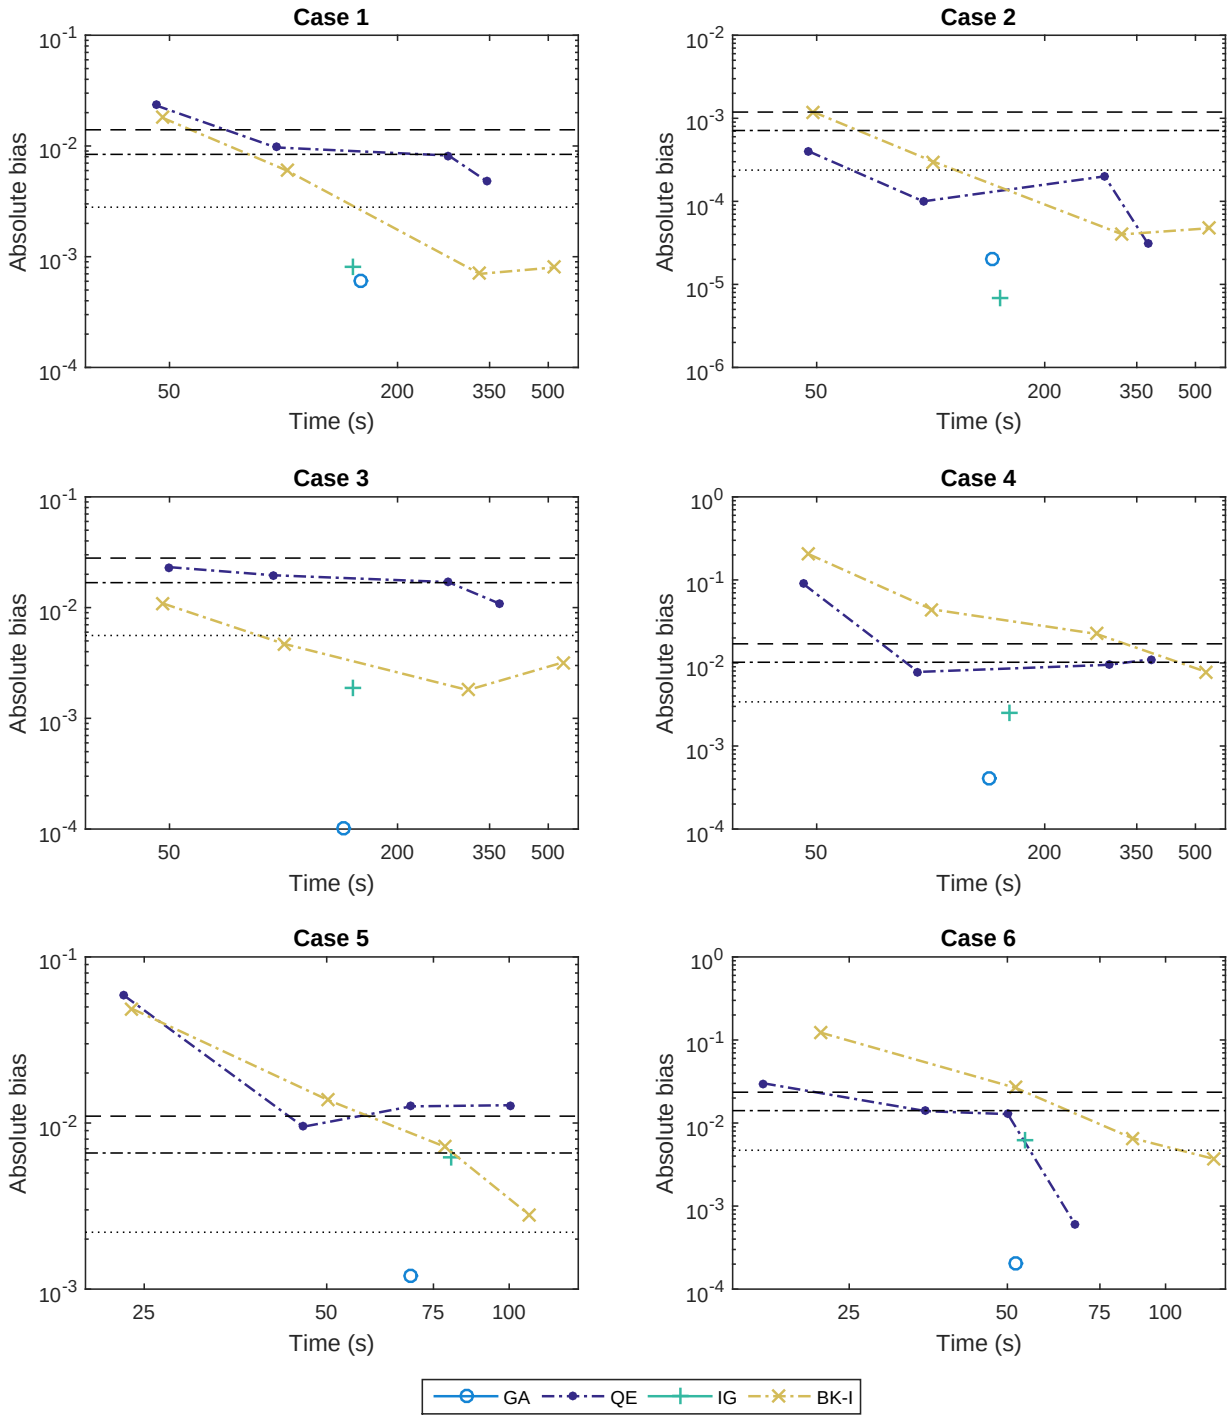

Figure 10: Absolute bias against computational time on pricing Asian options. The light blue line (circle marker) is used for the GA scheme, navy blue (dot marker) for QE, yellow (cross marker) for BK-I, and turquoise (plus marker) for IG. Standard deviation of the price (dotted line), 3 times the standard deviation (dashed-dot line), and 5 times the stardard deviation (large-dash line) are included.

Table 3: Results for the European call option using Cases 2 and 3.

Panel A: Case 2, European call option

QE			GA			IG			BK-I		
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
1.6	169.690	18.79	0.4	117.691	17.94	0.6	33.708	21.99	0.8	112.856	18.54
2	111.098	23.07	0.5	64.779	22.34	0.7	29.820	25.86	1	87.770	23.72
2.4	4.564	28.13	0.6	**0.879**	27.03	0.8	2.231	29.71	1.2	5.917	28.57
2.8	3.009	33.24	0.7	**0.575**	31.55	0.9	2.198	33.26	1.4	4.497	33.34
3.2	1.657	37.77	0.8	1.014	35.96	1	1.082	36.97	1.6	3.212	38.40
3.6	**0.710**	42.14	0.9	1.116	40.61	1.1	1.285	41.06	1.8	3.516	43.63
4	**0.811**	46.94	1	1.352	45.02	1.2	**0.642**	44.64	2	2.637	48.60
4.4	**0.575**	51.27	1.1	**0.676**	49.42	1.3	**0.575**	48.23	2.2	1.555	53.50
4.8	0.913	55.55	1.2	**0.710**	53.99	1.4	**0.744**	52.07	2.4	1.454	58.61
5.2	**0.744**	63.30	1.3	**0.642**	59.74	1.5	**0.473**	55.88	2.6	1.082	63.84
5.6	**0.135**	65.08	1.4	**0.879**	62.31	1.6	**0.135**	59.25	2.8	**0.879**	68.35
6	**0.541**	69.47	1.5	**0.879**	66.85	1.7	**0.068**	63.10	3	1.386	73.37
6.4	**0.068**	74.30	1.6	**0.473**	71.39	1.8	**0.744**	66.90	3.2	1.048	78.67
6.8	**0.237**	79.03	1.7	**0.642**	75.88	1.9	**0.034**	70.86	3.4	**0.845**	83.58
7.2	**0.169**	83.73	1.8	**0.034**	80.49	2	**0.473**	74.55	3.6	0.981	89.25
7.6	**0.068**	88.63	1.9	**0.068**	85.01	2.1	**0.169**	78.52	3.8	**0.237**	94.47
8	**0.710**	93.08	2	**0.068**	89.48	2.2	**0.034**	82.43	4	**0.541**	99.80

Panel B: Case 3, European call option

QE			GA			IG			BK-I		
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
1.6	0.830	18.67	0.4	0.214	17.96	0.6	0.240	21.96	0.8	1.155	18.59
2	0.429	23.41	0.5	**0.005**	22.42	0.7	0.171	25.56	1	0.808	23.31
2.4	0.246	28.85	0.6	**0.053**	27.15	0.8	**0.046**	29.65	1.2	0.552	28.59
2.8	0.122	33.05	0.7	**0.037**	31.69	0.9	**0.047**	33.28	1.4	0.388	33.64
3.2	**0.057**	38.11	0.8	0.072	36.31	1	**0.025**	37.28	1.6	0.336	38.86
3.6	**0.023**	43.05	0.9	**0.003**	40.87	1.1	**0.013**	40.86	1.8	0.267	43.65
4	**0.023**	47.19	1	**0.003**	45.04	1.2	**0.022**	44.70	2	0.210	48.64
4.4	**0.023**	51.64	1.1	**0.016**	49.57	1.3	**0.013**	48.65	2.2	0.177	53.52
4.8	**0.033**	55.95	1.2	**0.041**	54.05	1.4	**0.015**	52.23	2.4	0.181	58.69
5.2	**0.038**	60.93	1.3	**0.015**	58.80	1.5	**0.041**	56.07	2.6	0.111	63.98
5.6	**0.032**	65.19	1.4	**0.012**	62.44	1.6	**0.023**	59.31	2.8	0.093	68.34
6	**0.021**	69.78	1.5	**0.006**	66.91	1.7	**0.009**	63.24	3	0.090	73.60
6.4	0.071	74.47	1.6	**0.003**	71.39	1.8	**0.013**	67.03	3.2	0.086	78.82
6.8	**0.005**	79.35	1.7	**0.018**	75.94	1.9	**0.015**	70.98	3.4	**0.064**	83.90
7.2	**0.050**	84.23	1.8	**0.001**	80.51	2	**0.025**	74.67	3.6	**0.052**	89.23
7.6	**0.035**	88.48	1.9	**0.001**	85.07	2.1	**0.031**	78.66	3.8	**0.052**	94.57
8	**0.006**	93.20	2	**0.013**	89.66	2.2	**0.013**	82.48	4	0.066	99.90

fit and could approximate the true behaviour of the integrated variance over time.

Some extensions could be made in that direction, but whether this could be achieved in a computationally competitive algorithm is still an open question.

A European and Asian options, Cases 2 to 6

The relative biases for which their absolute counterparts are not statistically different from zero are in bold in Tables 3, 4 and 5.

Table 4: Results for the European call option using Cases 4, 5 and 6.

Panel A: Case 4, European call option

QE			GA			IG			BK-I		
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
1.6	9.453	19.03	0.4	2.655	18.62	0.6	1.744	22.98	0.8	14.868	18.48
2	5.934	23.45	0.5	1.404	22.94	0.7	1.560	26.61	1	10.580	23.51
2.4	3.970	27.90	0.6	0.947	27.34	0.8	1.301	30.46	1.2	7.707	27.94
2.8	2.657	32.52	0.7	0.209	31.91	0.9	1.053	34.33	1.4	5.838	32.85
3.2	1.868	37.31	0.8	**0.080**	36.42	1	0.791	38.31	1.6	4.581	37.78
3.6	1.314	41.95	0.9	0.158	40.92	1.1	0.556	42.10	1.8	3.596	42.77
4	0.966	46.40	1	0.194	45.44	1.2	0.329	46.08	2	2.958	47.82
4.4	0.602	51.07	1.1	0.133	49.96	1.3	0.294	50.01	2.2	2.452	52.99
4.8	0.529	55.60	1.2	0.137	54.55	1.4	0.192	53.92	2.4	2.057	58.11
5.2	0.342	60.30	1.3	**0.042**	59.10	1.5	**0.101**	57.89	2.6	1.732	63.26
5.6	0.231	65.12	1.4	**0.043**	63.77	1.6	**0.031**	61.94	2.8	1.566	68.55
6	0.210	69.71	1.5	**0.026**	68.34	1.7	**0.007**	66.01	3	1.314	73.86
6.4	**0.084**	74.33	1.6	**0.012**	73.54	1.8	**0.052**	72.76	3.2	1.216	79.17
6.8	**0.083**	79.13	1.7	**0.003**	77.50	1.9	**0.013**	74.14	3.4	1.049	84.56
7.2	**0.025**	83.79	1.8	**0.043**	82.10	2	**0.061**	78.12	3.6	0.903	89.88
7.6	**0.031**	88.56	1.9	**0.006**	86.80	2.1	**0.004**	82.16	3.8	0.835	95.36
8	**0.044**	93.41	2	**0.025**	91.72	2.2	**0.022**	86.22	4	0.793	100.79

Panel B: Case 5, European call option

QE			GA			IG			BK-I		
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
1.6	8.958	9.45	0.4	7.629	9.38	0.6	1.722	11.54	0.8	7.762	9.28
2.4	4.034	13.97	0.6	1.854	13.75	0.8	1.415	15.30	1.2	4.834	14.02
3.2	1.938	18.58	0.8	0.127	18.29	1	0.944	19.14	1.6	2.999	18.94
4	0.875	23.21	1	0.328	22.78	1.2	0.599	23.07	2	1.981	23.94
4.8	0.392	27.97	1.2	0.342	27.31	1.4	0.368	27.02	2.4	1.413	29.08
5.6	0.220	32.40	1.4	0.199	31.88	1.6	0.168	30.98	2.8	1.011	34.29
6.4	**0.009**	37.31	1.6	**0.077**	36.53	1.8	**0.087**	35.07	3.2	0.848	39.65
7.2	0.098	41.93	1.8	**0.086**	41.09	2	**0.002**	39.12	3.6	0.601	45.05
8	**0.045**	46.47	2	**0.042**	45.70	2.2	**0.038**	43.16	4	0.540	50.46
8.8	0.102	51.00	2.2	**0.009**	50.35	2.4	**0.041**	47.35	4.4	0.404	56.03
9.6	0.110	55.82	2.4	**0.016**	55.03	2.6	**0.009**	51.38	4.8	0.371	61.63
10.4	**0.084**	60.45	2.6	**0.036**	59.65	2.8	**0.011**	55.52	5.2	0.307	67.34
11.2	0.108	65.14	2.8	**0.036**	64.36	3	**0.002**	59.58	5.6	0.290	73.02
12	0.125	69.46	3	**0.029**	69.06	3.2	**0.046**	63.81	6	0.211	78.75

Panel C: Case 6, European call option

QE			GA			IG			BK-I		
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
2	1.770	8.92	0.5	0.394	7.31	0.25	0.267	7.65	1	4.668	10.00
3	0.752	13.25	0.75	**0.133**	11.30	0.5	0.191	11.71	1.5	2.188	15.85
4	0.466	17.26	1	**0.027**	15.25	0.75	0.153	11.78	2	1.206	22.31
5	0.264	21.53	1.25	**0.046**	19.46	1	**0.123**	15.93	2.5	0.804	29.23
6	0.189	26.21	1.5	**0.015**	23.92	1.25	**0.034**	20.23	3	0.607	36.52
7	**0.101**	30.08	1.75	**0.019**	27.97	1.5	**0.025**	24.86	3.5	0.317	43.88
8	**0.017**	34.14	2	**0.115**	32.60	1.75	**0.119**	34.22	4	0.291	51.82
9	**0.043**	38.25	2.25	**0.015**	37.28	2	**0.003**	38.99	4.5	0.198	60.14
10	**0.129**	42.41	2.5	**0.067**	42.01	2.25	**0.031**	44.05	5	**0.088**	68.64
11	**0.080**	46.70	2.75	**0.046**	47.02	2.5	**0.133**	49.25	5.5	0.220	77.40
12	**0.048**	50.82	3	**0.034**	51.97	2.75	**0.009**	54.30	6	**0.131**	86.38

Table 5: Results for the Asian call option using Cases 2 to 6.

Panel A: Case 2, Asian call option											
	QE			GA			IG			BK-I	
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
4	**1.645**	47.63	1	**1.645**	45.09	1	**2.878**	37.35	2	4.934	48.82
8	**0.411**	95.59	2	**0.822**	90.61	2	**1.645**	76.05	4	**1.234**	101.83
12	**0.822**	287.94	3	**0.082**	145.39	3	**0.411**	116.82	6	**0.165**	320.60
16	**0.130**	373.66	4	**0.057**	228.58	4	**0.028**	152.57	8	**0.195**	541.43
Panel B: Case 3, Asian call option											
	QE			GA			IG			BK-I	
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
4	0.071	49.89	1	**0.010**	49.58	1	**0.042**	40.82	2	**0.033**	48.04
8	0.060	93.86	2	**0.018**	89.56	2	**0.025**	74.43	4	**0.014**	100.38
12	0.052	272.05	3	**0.000**	144.20	3	**0.016**	113.97	6	**0.006**	306.57
16	**0.033**	373.24	4	**0.036**	223.69	4	**0.006**	153.08	8	**0.010**	547.00
Panel C: Case 4, Asian call option											
	QE			GA			IG			BK-I	
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
4	0.895	46.11	1	0.318	45.36	1	0.938	38.05	2	2.035	47.72
8	**0.078**	92.65	2	**0.042**	90.82	2	**0.015**	77.45	4	0.443	100.52
12	**0.095**	295.40	3	**0.004**	143.12	3	**0.007**	118.58	6	0.224	273.60
16	0.111	383.22	4	**0.030**	211.19	4	**0.025**	160.99	8	**0.078**	534.27
Panel D: Case 5, Asian call option											
	QE			GA			IG			BK-I	
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
4	0.806	23.15	1	0.142	22.79	1	1.164	19.12	2	0.670	23.88
8	0.131	45.73	2	0.098	0.10	2	**0.043**	38.90	4	0.189	50.37
12	0.173	68.93	3	**0.017**	68.80	3	**0.047**	59.28	6	0.099	78.33
16	0.176	100.49	4	**0.034**	92.47	4	**0.085**	80.36	8	**0.039**	107.93
Panel E: Case 6, Asian call option											
	QE			GA			IG			BK-I	
m	Bias	Time	m	Bias	Time	m	Bias	Time	m	Bias	Time
4	0.310	17.16	1	**0.036**	15.02	1	**0.044**	15.70	2	1.270	22.00
8	**0.143**	35.00	2	**0.063**	33.90	2	**0.037**	34.84	4	0.274	51.74
12	**0.132**	50.09	3	**0.002**	51.82	3	**0.063**	54.05	6	**0.068**	86.29
16	**0.006**	67.29	4	**0.004**	72.24	4	**0.004**	75.27	8	**0.038**	123,53

References

[1] Ahrens, J., and U. Dieter. 1980. Sampling from Binomial and Poisson distributions: A method with bounded computation times. *Computing* 25:193–208.

[2] Andersen, L. 2008. Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance* 11:1–42.

[3] Andersen, L., and V. V. Piterbarg. 2007. Moment explosions in stochastic volatility models. *Finance and Stochastics* 11:29–50.

[4] Bakshi, G., C. Cao, and Z. Chen. 1997. Empirical performance of alternative option pricing models. *Journal of Finance* pp. 2003–2049.

[5] Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–654.

[6] Broadie, M., and O. Kaya. 2004. Exact simulation of option greeks under stochastic volatility and jump diffusion models. *Proceedings of the 2004 Winter Simulation Conference* pp. 1607–1615.

[7] Broadie, M., and O. Kaya. 2006. Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* 54:217–231.

[8] Cox, J., J. Ingersoll Jr, and S. Ross. 1985. A theory of the term structure of interest rates. *Econometrica* 53:385–408.

[9] Devroye, L. 1986. *Non-uniform random variate generation*. Springer-Verlag.

[10] Dufresne, D. 2001. The integrated square-root process. *Research paper of the Centre for Actuarial Studies* 90:1–34.

[11] Feller, W. 1951. Two singular diffusion problems. *Annals of Mathematics* 54:173–182.

[12] Fishman, G. 1996. *Monte Carlo: Concepts, algorithms, and applications*. Springer.

[13] Gatheral, J. 2006. *The volatility surface: a practitioner's guide*. John Wiley.

[14] Gil-Pelaez, J. 1951. Note on the inversion theorem. *Biometrika* 38:481–482.

[15] Glasserman, P., and K. Kim. 2011. Gamma expansion of the Heston stochastic volatility model. *Finance and Stochastics* 15:267–296.

[16] Heston, S. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6:327–343.

[17] Hull, J., and A. White. 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42:281–300.

[18] Johnson, N., S. Kotz, and N. Balakrishnan. 1995. *Continuous univariate distributions*. John Wiley.

[19] Kahl, C., and P. Jäckel. 2006. Fast strong approximation Monte Carlo schemes for stochastic volatility models. *Quantitative Finance* 6:513–536.

[20] Kotz, S., N. Johnson, and N. Balakrishnan. 2000. *Continuous multivariate distributions: models and applications*. John Wiley.

[21] Lord, R., R. Koekkoek, and D. van Dijk. 2010. A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance* 10:177–194.

[22] Marsaglia, G., and W. Tsang. 2000. A simple method for generating gamma variables. *ACM Transactions on Mathematical Software* 26:363–372.

[23] Merton, R. 1973. Theory of rational option pricing. *Bell Journal of Economics* 4:141–183.

[24] Michael, J., W. Schucany, and R. Haas. 1976. Generating random variates using transformations with multiple roots. *American Statistician* 30:88–90.

[25] Milstein, G. 1978. A method of second order accuracy integration of stochastic differential equations. *Teoriya Veroyatnostei i ee Primeneniya* 23:414–419.

[26] Moodley, N. 2005. The Heston model: A practical approach with Matlab code. Working paper.

[27] Patnaik, P. 1949. The non-central χ^2-and F-distribution and their applications. *Biometrika* 36:202–232.

[28] Smith, R. 2007. An almost exact simulation method for the Heston model. *Journal of Computational Finance* 11:115.

[29] Stewart, T., L. Strijbosch, H. Moors, and P. Batenburg. 2007. A simple approximation to the convolution of gamma distributions. Working paper.

[30] Tse, S., and J. Wan. 2013. Low-bias simulation scheme for the Heston model by Inverse Gaussian approximation. *Quantitative Finance* 13:919–937.

[31] van Haastrecht, A., and A. Pelsser. 2010. Efficient, almost exact simulation of the Heston stochastic dvolatility model. *International Journal of Theoretical and Applied Finance* 13:1–43.

[32] Zhu, J. 2011. A simple and exact simulation approach to Heston model 18:26–36. Journal of Derivatives.