



Research Paper

A mixed Monte Carlo and partial differential equation variance reduction method for foreign exchange options under the Heston–Cox–Ingersoll–Ross model

Andrei Cozma and Christoph Reisinger

Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK;
emails: andrei.cozma@maths.ox.ac.uk, christoph.reisinger@maths.ox.ac.uk

(Received September 2, 2015; revised January 13, 2016; accepted March 21, 2016)

ABSTRACT

In this paper, we consider the valuation of European and path-dependent options in foreign exchange markets when the currency exchange rate evolves according to the Heston model combined with the Cox–Ingersoll–Ross (CIR) dynamics for the stochastic domestic and foreign short interest rates. The mixed Monte Carlo/partial differential equation method requires that we simulate only the paths of the squared volatility and the two interest rates, while an “inner” Black–Scholes-type expectation is evaluated by means of a partial differential equation. This can lead to a substantial variance reduction and complexity improvements under certain circumstances depending on the contract and the model parameters. In this work, we establish the uniform boundedness of moments of the exchange rate process and its approximation, and prove strong convergence of the latter in L^p ($p \geq 1$). Then, we carry out a variance reduction analysis and obtain accurate approximations for quantities of interest. All theoretical contributions can be extended to multi-factor short rates in a straightforward manner. Finally, we illustrate the efficiency of the method for the four-factor Heston–CIR model through a detailed quantitative assessment.

Keywords: conditional Monte Carlo; mixed Monte Carlo/PDE; stochastic volatility; stochastic interest rates; variance reduction; strong convergence.

1 INTRODUCTION

In foreign exchange (FX) markets, option pricing with stochastic volatility and stochastic interest rates has seen a large amount of interest in recent years (Ahlip and Rutkowski 2013; Grzelak and Oosterlee 2011; Van Haastrecht and Pelsser 2011), leading to the extension of the Heston (1993) two-factor stochastic volatility model and the Schöbel–Zhu model (Schöbel and Zhu 1999) to currency derivatives. Although appealing due to its simplicity, assuming constant interest rates is inappropriate for long-dated FX products, and the effect of interest rate volatility can even outweigh that of FX rate volatility for long maturities, as has been confirmed by empirical results (Van Haastrecht *et al* 2009). Here, the spot FX rate is defined as the number of units of domestic currency per unit of foreign currency.

In this paper, we consider the four-factor Heston–Cox–Ingersoll–Ross (Heston–CIR) model proposed and examined in Ahlip and Rutkowski (2013) when the volatility and the exchange rate dynamics are correlated, whereas the domestic and foreign interest rates are pairwise independent and also independent of the exchange rate and the volatility. Our motivation comes from the fact that the square-root (CIR) process (Cox *et al* 1985) for the variance and interest rates is widely used in the industry due to its desirable properties, such as mean-reversion and nonnegativity. Under these restrictive assumptions on the independence of the Brownian drivers, Cox *et al* argue that the model is affine and derive a semi-analytical formula for the European call option price. The importance of a nonzero correlation between the FX rate and the interest rate(s) is recognized in Hunter (2005). However, any other nonzero correlations give rise to a non-affine model, in which case we lose analytical tractability.

Therefore, we find ourselves forced to turn to numerical algorithms, and the mixed Monte Carlo/partial differential equation (PDE) solver (Loeper and Pironneau 2009) is a good alternative to the classical Monte Carlo and finite-difference (FD) methods. Monte Carlo simulation methods (Glasserman 2003) can handle path-dependent features easily and scale linearly with the dimension. However, a considerable number of simulations are typically required for a good accuracy. Conversely, finite-difference methods easily incorporate early exercise features and provide a fast convergence for low-dimensional problems (up to three dimensions), but become intractable as the dimensionality increases. This makes Monte Carlo methods more attractive for the four-factor Heston–CIR model, and the relatively slow convergence $\mathcal{O}(M^{-0.5})$ in

the number of simulations M raises the question of how to find an efficient variance reduction technique.

The idea behind the mixed Monte Carlo/PDE method is to write the option values as nested conditional expectations. Then, the innermost expectation is evaluated analytically, in the case of European-style options, or using finite differences, and the outer expectation evaluated by simulation. The earliest related published work is by Hull and White (1987), who consider a European call option under a two-dimensional stochastic volatility model with uncorrelated asset price and volatility and prove that, conditional on the integral of the variance process, the asset price is lognormally distributed, and hence the option price can be expressed as a Black–Scholes price. Willard (1997) extends the analysis of Hull and White to path-independent options under stochastic volatility and instantaneously correlated factors, and uses the smoothness of the “conditional price” to calculate price sensitivities (the Greeks). Willard also employs quasi-Monte Carlo (low-discrepancy) methods to further reduce the variance of price estimates, but with no effect on the discretization bias. The mixed Monte Carlo/PDE method develops this conditioning technique, known as conditional Monte Carlo, and allows the combined use of Monte Carlo and finite-difference methods for the valuation of path-dependent contracts.

The utility of the method can be easily recognized when pricing European-style options under the Heston–CIR model. Conditional on the entire paths of the squared volatility and the domestic and foreign interest rates, the dynamics of the exchange rate are governed by a geometric Brownian motion with time-dependent drift and diffusion coefficients. Combined with the existence of a closed-form solution for the conditional option price, the algorithm reduces the variance in Monte Carlo simulations – by eliminating a source of noise – and the dimension of the problem, from 4 to 3. The mixed Monte Carlo/PDE method has been the subject of several numerical studies in the past few years (see, for example, Ang 2013; Dang *et al* 2015; Lipp *et al* 2013; McGhee 2014), and various extensions to the original idea have been considered. For instance, Dang *et al* use the Hull–White-type dynamics of the interest rates and condition on the variance path to find a closed-form solution for the conditional price of a European option. A few of these works examine convergence properties of the algorithm using heuristic arguments, but none of them take into account the error arising from the discretization of the variance and interest rate processes.

To the best of our knowledge, the convergence of the mixed Monte Carlo/PDE method has not yet been established, and even the literature on Monte Carlo methods under stochastic volatility is scarce. Higham and Mao (2005) considered an Euler simulation of the Heston model with a reflection fix and proved strong convergence of the stopped approximation process and weak convergence for a European put and an up-and-out call, by using the boundedness of payoffs. Cozma and Reisinger

(2015a) extended these results to derivatives with unbounded payoffs, stochastic–local volatility models, stochastic interest rates and exotic payoffs. Their discretization scheme coincides with that considered in this paper at the discrete time points, but necessarily differs in the continuous-time interpolation. Some results from the earlier work can be utilized, although stronger conditions on the model parameters were needed there and therefore the results here can also be seen as an improvement for the standard (ie, non-mixed) scheme. In particular, the new results are always guaranteed for zero or negative correlation in the Heston model. The main conceptual difference, however, is the new continuous-time interpolation using conditional drifts, which is crucial for deriving the conditional PDEs and leads to different technical challenges. Also, for path-dependent options, neither the original scheme nor the analysis applies to the present work.

In this paper, we study convergence properties of the mixed Monte Carlo/PDE method with the full truncation Euler (FTE) discretization for the squared volatility and the two interest rates, and demonstrate the efficiency of the method for a European call and an up-and-out put option. We prefer the full truncation scheme (Lord *et al* 2010) because it preserves positivity, is easy to implement and is found empirically to produce the smallest bias of all Euler schemes. An interesting alternative to the FTE scheme would be the backward Euler–Maruyama (BEM) scheme (Neuenkirch and Szpruch 2014). However, the quanto correction term in the dynamics of the foreign interest rate would lead to technical challenges in the convergence analysis. The major contributions of this paper are as follows.

- We establish the uniform boundedness of moments of the four-dimensional process and its approximation, and prove strong convergence of the discretization scheme in L^p ($p \geq 1$). Then, we deduce the convergence of mixed Monte Carlo/PDE estimators for computing option prices and discuss possible extensions to higher-dimensional models.
- We carry out a thorough theoretical variance reduction analysis of the mixed Monte Carlo/PDE method and employ standard Monte Carlo – with the log-Euler discretization – as the reference method, noting that the analysis applies to general interest rate dynamics. In particular, we investigate how different values of the underlying model parameters affect the variance of mixed estimators.
- We perform a series of numerical experiments and demonstrate the convergence of the mixed Monte Carlo/PDE method under the four-factor FX model for two financial derivatives: a European call and an up-and-out put option. In addition, we establish the efficiency of the method by comparison with alternative numerical schemes and examine the sensitivity of the variance reduction factor to changes in the parameters.

The remainder of this paper is structured as follows. In Section 2 we introduce the four-factor FX model, define the mixed simulation scheme and describe the pricing algorithm. In Section 3 we prove strong convergence of the exchange rate approximations and then discuss some extensions. Detailed proofs of some technical results are given in the online appendix. In Section 4 we carry out a variance reduction analysis of the mixed Monte Carlo/PDE method for a European option. Various numerical experiments are presented and discussed in Section 5. Finally, Section 6 summarizes the results and outlines possible future work.

2 PRELIMINARIES

2.1 The four-factor model

We have in mind a model in an FX market, for the spot FX rate S , the variance of the FX rate v , the domestic short interest rate r^d and the foreign short interest rate r^f . Unless otherwise stated, the subscripts and superscripts “d” and “f” in this paper are used to indicate “domestic” and “foreign”, respectively. Consider a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{Q})$ and suppose that the dynamics of the underlying processes are governed by the following system of stochastic differential equations (SDEs) under the domestic risk-neutral measure \mathbb{Q} :

$$\left. \begin{aligned} dS_t &= (r_t^d - r_t^f)S_t dt + \sqrt{v_t}S_t dW_t^s, \\ dv_t &= k(\theta - v_t) dt + \xi \sqrt{v_t} dW_t^v, \\ dr_t^d &= k_d(\theta_d - r_t^d) dt + \xi_d \sqrt{r_t^d} dW_t^d, \\ dr_t^f &= (k_f \theta_f - k_f r_t^f - \rho_{sf} \xi_f \sqrt{v_t r_t^f}) dt + \xi_f \sqrt{r_t^f} dW_t^f, \end{aligned} \right\} \quad (2.1)$$

where $\{W^s, W^v, W^d, W^f\}$ are correlated standard Brownian motions (BM) under the risk-neutral measure with constant correlation matrix Σ . We consider a full correlation structure between the four Brownian drivers, which reflects movements in the financial markets more accurately and allows for better calibrations. Then, we decouple $\{W^s, W^f, W^d, W^v\}$ and express them as linear combinations of independent Brownian motions $\{W^1, W^2, W^3, W^4\}$. Hence, we define the two vectors $W = [W^s, W^f, W^d, W^v]^T$ and $\tilde{W} = [W^1, W^2, W^3, W^4]^T$. As Σ is symmetric positive definite, a standard Cholesky factorization gives rise to an upper triangular matrix of coefficients $A = (a_{ij})_{1 \leq i, j \leq 4}$, satisfying $\Sigma = AA^T$:

$$\Sigma = \begin{bmatrix} 1 & \rho_{sf} & \rho_{sd} & \rho_{sv} \\ \rho_{sf} & 1 & \rho_{df} & \rho_{vf} \\ \rho_{sd} & \rho_{df} & 1 & \rho_{vd} \\ \rho_{sv} & \rho_{vf} & \rho_{vd} & 1 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ 0 & a_{22} & a_{23} & a_{24} \\ 0 & 0 & a_{33} & a_{34} \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (2.2)$$

This decomposition implies that we can choose \tilde{W} such that $W = A\tilde{W}$. We can determine the matrix of coefficients by solving a system of ten equations. Assuming $\rho_{vd} \neq \pm 1$, we find

$$\begin{aligned} a_{14} &= \rho_{sv}, & a_{24} &= \rho_{vf}, \\ a_{34} &= \rho_{vd}, & a_{33} &= (1 - \rho_{vd}^2)^{1/2}, \\ a_{13} &= (\rho_{sd} - \rho_{sv}\rho_{vd})(1 - \rho_{vd}^2)^{-1/2}, \\ a_{23} &= (\rho_{df} - \rho_{vf}\rho_{vd})(1 - \rho_{vd}^2)^{-1/2}, \\ a_{22} &= (1 - \rho_{df}^2 - \rho_{vf}^2 - \rho_{vd}^2 + 2\rho_{df}\rho_{vf}\rho_{vd})^{1/2}(1 - \rho_{vd}^2)^{-1/2}. \end{aligned}$$

a_{11} and a_{12} can be found in a similar fashion.

The quanto correction term in the drift of the foreign interest rate in (2.1) comes from changing from the foreign risk-neutral measure to the domestic one (Clark 2011). Alternatively, we can also think of (2.1) as a model in an equity market with asset price process S , interest rate r^d and dividend yield r^f , in which case the quanto drift-adjustment term vanishes.

2.2 The mixed simulation scheme

First, we discretize the variance and the two interest rate processes using the FTE scheme of Lord *et al* (2010). Consider the square-root process

$$dy_t = k_y(\theta_y - y_t)dt + \xi_y \sqrt{y_t} dW_t^y. \quad (2.3)$$

For a time interval $[0, T]$, consider a uniform grid: $\delta t = T/N$, $t_n = n\delta t$ for all $n \in \{0, 1, \dots, N\}$. We introduce the discrete-time auxiliary process

$$\tilde{y}_{t_{n+1}} = \tilde{y}_{t_n} + k_y(\theta_y - \tilde{y}_{t_n}^+)\delta t + \xi_y \sqrt{\tilde{y}_{t_n}^+} \delta W_{t_n}^y, \quad (2.4)$$

where $y^+ = \max(0, y)$ and $\delta W_{t_n}^y = W_{t_{n+1}}^y - W_{t_n}^y$, and the continuous-time interpolation

$$\tilde{y}_t = \tilde{y}_{t_n} + k_y(\theta_y - \tilde{y}_{t_n}^+)(t - t_n) + \xi_y \sqrt{\tilde{y}_{t_n}^+}(W_t^y - W_{t_n}^y), \quad \forall t \in [t_n, t_{n+1}), \quad (2.5)$$

as suggested in Higham and Mao (2005). Moreover, we define the nonnegative processes

$$Y_t = \tilde{y}_t^+ \quad (2.6)$$

and

$$\bar{Y}_t = \tilde{y}_{t_n}^+, \quad (2.7)$$

whenever $t \in [t_n, t_{n+1})$. Let \bar{V} and \bar{r}^d be the FTE discretizations – as defined in (2.7) – of the variance and the domestic interest rate, respectively. Taking into account the presence of the quanto correction term in the drift of the foreign interest rate, we similarly define

$$\tilde{r}_t^f = \tilde{r}_{t_n}^f + \left[k_f \theta_f - k_f (\tilde{r}_{t_n}^f)^+ - \rho_{sf} \xi_f \sqrt{\tilde{v}_{t_n}^+ (\tilde{r}_{t_n}^f)^+} \right] (t - t_n) + \xi_f \sqrt{(\tilde{r}_{t_n}^f)^+} (W_t^f - W_{t_n}^f), \quad (2.8)$$

as well as

$$\hat{r}_t^f = (\tilde{r}_t^f)^+ \quad (2.9)$$

and

$$\bar{r}_t^f = (\tilde{r}_{t_n}^f)^+, \quad (2.10)$$

whenever $t \in [t_n, t_{n+1})$. Next, we define \bar{S} , the continuous-time approximation of S , as the solution to the following SDE:

$$\left. \begin{aligned} d\bar{S}_t &= \mu_t \bar{S}_t dt + a_{11} \sqrt{\bar{V}_t} \bar{S}_t dW_t^1, \\ \mu_t &= \bar{r}_t^d - \bar{r}_t^f - \frac{1}{2}(1 - a_{11}^2) \bar{V}_t + \sum_{j=2}^4 a_{1j} \sqrt{\bar{V}_t} \frac{\delta W_t^j}{\delta t}, \end{aligned} \right\} \quad (2.11)$$

where $\delta W_t^j = W_{t_{n+1}}^j - W_{t_n}^j$ for all $t \in [t_n, t_{n+1})$. Hence, μ_t is piecewise constant. For convenience, we introduce the actual and the approximated log-processes, $x = \log S$ and $X = \log \bar{S}$.

Conditioning on the trajectories of $\{W^j, j = 2, 3, 4\}$, ie, on the complete knowledge of the paths of the variance and interest rates, \bar{S} evolves like a geometric Brownian motion with time-dependent drift and diffusion coefficients. It follows from Ito's formula that

$$\begin{aligned} \bar{S}_T &= S_0 \exp \left\{ \int_0^T (\bar{r}_u^d - \bar{r}_u^f - \frac{1}{2} \bar{V}_u) du \right. \\ &\quad \left. + \sum_{j=2}^4 a_{1j} \int_0^T \sqrt{\bar{V}_u} \frac{\delta W_u^j}{\delta t} du + a_{11} \int_0^T \sqrt{\bar{V}_u} dW_u^1 \right\}. \end{aligned}$$

However, we know the conditional probability law of the stochastic integral on the right-hand side to be that of a normal random variable, namely

$$\int_0^T \sqrt{\bar{V}_u} dW_u^1 \stackrel{\text{law}}{=} \sqrt{\int_0^T \bar{V}_u du} Z,$$

where $Z \sim \mathcal{N}(0, 1)$, so we can think of \bar{S}_T as a function of Z . Therefore, we can express it as

$$\bar{S}_T \stackrel{\text{law}}{=} S_0 \exp\{(r - q - \tfrac{1}{2}\sigma^2)T + \sigma\sqrt{T}Z\}, \quad (2.12)$$

where

$$r = \frac{1}{T} \int_0^T \bar{r}_u^d du = \frac{1}{N} \sum_{i=0}^{N-1} \bar{r}_{t_i}^d, \quad \sigma^2 = \frac{a_{11}^2}{T} \int_0^T \bar{V}_u du = \frac{a_{11}^2}{N} \sum_{i=0}^{N-1} \bar{V}_{t_i}$$

and

$$\begin{aligned} q &= \frac{1}{T} \int_0^T \bar{r}_u^f du + \frac{1 - a_{11}^2}{2T} \int_0^T \bar{V}_u du - \frac{1}{T} \sum_{j=2}^4 a_{1j} \int_0^T \sqrt{\bar{V}_u} \frac{\delta W_u^j}{\delta t} du \\ &= \frac{1}{N} \sum_{i=0}^{N-1} \bar{r}_{t_i}^f + \frac{1 - a_{11}^2}{2N} \sum_{i=0}^{N-1} \bar{V}_{t_i} - \frac{1}{T} \sum_{j=2}^4 a_{1j} \sum_{i=0}^{N-1} \sqrt{\bar{V}_{t_i}} \delta W_{t_i}^j. \end{aligned}$$

Conditional on the trajectories of $\{W^j, j = 2, 3, 4\}$, (2.12) has the same law as a terminal asset price that evolves as a geometric Brownian motion with interest rate r , continuous dividend yield q and volatility σ , all constant. Then the arbitrage-free price at time $t = 0$ of a European-style option with payoff $f(S_T)$ is the discounted expectation under the risk-neutral measure, which can be approximated, using the mixed Monte Carlo/PDE method, by

$$\bar{U} = \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_u^d du \right) f(\bar{S}_T) \right] = \mathbb{E}[\mathbb{E}[e^{-rT} f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}]], \quad (2.13)$$

where $\{\mathcal{G}_t^{f,d,v}, 0 \leq t \leq T\}$ is the natural filtration generated by the independent Brownian motions $\{W^2, W^3, W^4\}$, ie, generated by the processes v, r^d, r^f as observed until time T . The second equality in (2.13) comes from the “tower property” of conditional expectations. Unless otherwise stated, all expectations are under \mathbb{Q} . Let the approximate conditional option price be the inner expectation in (2.13), which is analytically tractable for European contracts,

$$\mathbb{E}[e^{-rT} f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}] = e^{-rT} \int_{-\infty}^{\infty} f(\bar{S}_T(z)) \phi(z) dz,$$

where ϕ and Φ are the standard normal probability distribution function (PDF) and cumulative distribution function (CDF), respectively. The conditional prices of some popular financial instruments are given below:

$$\left. \begin{array}{ll} \text{European options:} & \psi S_0 e^{-qT} \Phi(\psi d_1) - \psi K e^{-rT} \Phi(\psi d_2), \\ \text{cash-or-nothing options:} & e^{-rT} \Phi(\psi d_2), \\ \text{asset-or-nothing options:} & S_0 e^{-qT} \Phi(\psi d_1), \end{array} \right\} \quad (2.14)$$

where $\psi = 1$ for a call and $\psi = -1$ for a put, and

$$d_{1,2} = \frac{\log(S_0/K) + (r - q \pm \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}. \quad (2.15)$$

The approximate option price, ie, the outer expectation in (2.13), is estimated by a Monte Carlo average over a sufficiently large number of discrete trajectories of $\{W^2, W^3, W^4\}$.

There are many other derivatives that admit a closed-form solution for the conditional price, eg, the power option, the chooser option or the forward-start option. However, for most path-dependent derivatives we need to use a different approach in order to compute the conditional price, in which case we will rely on finite-difference methods (see Section 5).

We conclude this section with a discussion on our choice of conditioning. For European option pricing, an analytical formula for the inner expectation is available only when conditioning on all three factors, which results in a reduction of the problem by one dimension. The high-dimensionality of the Heston–CIR model makes this the natural choice. For path-dependent option pricing, due to the quanto correction term in the drift of the foreign rate, v and r^f are coupled, and hence we cannot condition on r^f alone. Moreover, the variance of Monte Carlo estimators due to the short rates is typically much lower than the variance due to the instantaneous squared volatility. Hence, we could alternatively condition on r^d and solve a three-dimensional PDE for the inner expectation. On the one hand, we could reach the same level of accuracy with fewer Monte Carlo sample paths than when simulating v and r^f as well. On the other hand, the computational effort grows linearly with the dimension for Monte Carlo methods and exponentially for finite-difference methods. Hence, we believe that conditioning on all three factors is more efficient. However, if the exchange rate and the foreign interest rate dynamics are independent, the quanto correction term vanishes. In this case, conditioning on the two short rates would be an interesting alternative.

3 CONVERGENCE ANALYSIS

Even though weak convergence is very important in financial mathematics when estimating expectations of payoffs, strong convergence may be required for complex path-dependent derivatives and plays a key role in multilevel Monte Carlo methods (Giles 2008). In this section, we prove the strong convergence of the approximation scheme defined in (2.11). Hence, we examine first the exponential integrability properties of the square-root process and its discretization, and then the finiteness of moments of order higher than 1 of the exchange rate process and its approximation.

Let y be the square-root process defined in (2.3) and let \bar{Y} be the piecewise constant FTE interpolant from (2.7). The exponential integrability of functionals of the two processes was discussed in Cozma and Reisinger (2015b, Propositions 3.2 and 3.6). However, we need to adjust their second result for our approximation scheme (2.11) in order to establish the convergence.

LEMMA 3.1 *Let $\lambda, \mu \in \mathbb{R}$ be given, let $\Delta \equiv \lambda + \frac{1}{2}\mu^2$ and define the stochastic process*

$$\bar{\Theta}_t \equiv \exp \left\{ \lambda \int_0^t \bar{Y}_u du + \mu \int_0^t \sqrt{\bar{Y}_u} \frac{\delta W_u^y}{\delta t} du \right\}, \quad \forall t \in [0, T]. \quad (3.1)$$

If $\Delta \leq 0$ and $T \geq 0$, or if $\Delta > 0$ and $T \leq T^$, then there exists $\eta > 0$ such that*

$$\sup_{\delta t \in (0, \eta)} \sup_{t \in [0, T]} \mathbb{E}[\bar{\Theta}_t] < \infty, \quad (3.2)$$

where T^ is as given below.*

(1) *If $k_y \leq \xi_y(\mu + \sqrt{0.5\Delta})$,*

$$T^* = \frac{1}{\xi_y(\mu + \sqrt{2\Delta}) - k_y}. \quad (3.3)$$

(2) *If $k_y > \xi_y(\mu + \sqrt{0.5\Delta})$,*

$$T^* = \frac{2(k_y - \mu\xi_y)}{\xi_y^2 \Delta}. \quad (3.4)$$

PROOF See Appendix A (available online). □

3.1 Moment bounds

For many stochastic volatility models, moments of order higher than 1 can blow-up in finite time (Andersen and Piterbarg 2007). This can cause significant problems in practice, for instance, when computing the arbitrage-free price of an option whose payoff function has superlinear growth. The same troublesome behavior can be observed for the Euler–Maruyama approximation of some SDEs with superlinearly growing drift or diffusion coefficients, where moments diverge in finite time (Hutzenthaler and Jentzen 2015). Next, we prove the boundedness of moments of the exchange rate process and its approximation.

At this point, we assume that $\rho_{vd} \neq \pm 1$ and that a_{13} is nonzero, ie, $\rho_{sd} \neq \rho_{sv}\rho_{vd}$.

PROPOSITION 3.2 For $\alpha \geq 1$, define the two quantities

$$q_0(\alpha) \equiv \frac{1}{2\alpha^2\xi^2a_{13}^2} \left\{ \sqrt{[2\alpha\rho_{sv}\xi k + \alpha^2\xi^2(a_{11}^2 + a_{12}^2) - \alpha\xi^2]^2 + 4\alpha^2a_{13}^2\xi^2k^2} \right. \\ \left. - [2\alpha\rho_{sv}\xi k + \alpha^2\xi^2(a_{11}^2 + a_{12}^2) - \alpha\xi^2] \right\}, \quad (3.5)$$

$$q_1(\alpha) \equiv q_0(\alpha)\mathbf{1}_{\rho_{sv} \leq 0} + \min \left\{ q_0(\alpha), \frac{k}{\alpha\rho_{sv}\xi} \right\} \mathbf{1}_{\rho_{sv} > 0}. \quad (3.6)$$

If the conditions

$$k > \alpha\rho_{sv}\xi + \sqrt{\alpha(\alpha-1)}\xi \quad \text{and} \quad \frac{k_d^2}{2\xi_d^2} > \frac{\alpha q_1(\alpha)}{q_1(\alpha) - 1} \quad (3.7)$$

on the model parameters are satisfied, then there exists $\alpha_1 > \alpha$ such that, for all $\omega \in [1, \alpha_1)$,

$$\sup_{t \in [0, T]} \mathbb{E}[S_t^\omega] < \infty. \quad (3.8)$$

PROOF See Appendix B (available online). \square

Next, assume that a_{13} and a_{14} are not simultaneously zero, ie, $\rho_{sv}^2 + \rho_{sd}^2 \neq 0$.

PROPOSITION 3.3 For $\alpha \geq 1$, define

$$q_2(\alpha) \equiv \left\{ \sqrt{[\alpha\rho_{sv}\xi + \frac{1}{4}T\alpha^2\xi^2(a_{11}^2 + a_{12}^2) - \frac{1}{4}T\alpha\xi^2]^2 + T\alpha^2\xi^2(a_{13}^2 + a_{14}^2)k} \right. \\ \left. - [\alpha\rho_{sv}\xi + \frac{1}{4}T\alpha^2\xi^2(a_{11}^2 + a_{12}^2) - \frac{1}{4}T\alpha\xi^2] \right\} \\ \times \frac{2}{T\alpha^2\xi^2(a_{13}^2 + a_{14}^2)}. \quad (3.9)$$

If the conditions

$$k > \alpha\rho_{sv}\xi + \frac{1}{4}\alpha(\alpha-1)T\xi^2 \quad \text{and} \quad \frac{2k_d}{T\xi_d^2} > \frac{\alpha q_2(\alpha)}{q_2(\alpha) - 1} \quad (3.10)$$

on the model parameters are satisfied, then there exists $\alpha_2 > \alpha$ such that, for all $\omega \in [1, \alpha_2)$, we can find $\eta_\omega > 0$ such that

$$\sup_{\delta t \in (0, \eta_\omega)} \sup_{t \in [0, T]} \mathbb{E}[\bar{S}_t^\omega] < \infty. \quad (3.11)$$

PROOF See Appendix C (available online). \square

Since the most popular FX and equity contracts grow at most linearly in FX and asset prices, and their valuation requires the computation of the expected discounted payoff under the risk-neutral measure, it is useful to study finiteness of moments under discounting. Let R be the discounted exchange rate process,

$$R_t = S_0 \exp \left\{ - \int_0^t (r_u^f + \frac{1}{2} v_u) du + \int_0^t \sqrt{v_u} dW_u^s \right\}, \quad (3.12)$$

and let \bar{R} be its continuous-time approximation,

$$\begin{aligned} \bar{R}_t = S_0 \exp \left\{ - \int_0^t (\bar{r}_u^f + \frac{1}{2} \bar{V}_u) du \right. \\ \left. + a_{11} \int_0^t \sqrt{\bar{V}_u} dW_u^1 + \sum_{j=2}^4 a_{1j} \int_0^t \sqrt{\bar{V}_u} \frac{\delta W_u^j}{\delta t} du \right\}. \end{aligned} \quad (3.13)$$

PROPOSITION 3.4 *Let $\alpha \geq 1$. If $T < T^*$, there exists $\alpha_1 > \alpha$ such that, for all $\omega \in [1, \alpha_1)$,*

$$\sup_{t \in [0, T]} \mathbb{E}[R_t^\omega] < \infty. \quad (3.14)$$

If $\alpha > 1$ and $T \geq T^$, then*

$$\mathbb{E}[R_T^\alpha] = \infty. \quad (3.15)$$

If $\alpha = 1$, then $T^ = \infty$, whereas if $\alpha > 1$, then T^* is as given below.*

(1) *If $k < \alpha \rho_{sv} \xi - \sqrt{\alpha(\alpha-1)} \xi$,*

$$T^* = \frac{1}{v(\alpha)} \log \left(\frac{\alpha \rho_{sv} \xi - k + v(\alpha)}{\alpha \rho_{sv} \xi - k - v(\alpha)} \right), \quad (3.16)$$

where

$$v(\alpha) = \sqrt{(\alpha \rho_{sv} \xi - k)^2 - \alpha(\alpha-1) \xi^2}.$$

(2) *If $k = \alpha \rho_{sv} \xi - \sqrt{\alpha(\alpha-1)} \xi$,*

$$T^* = \frac{2}{\alpha \rho_{sv} \xi - k}. \quad (3.17)$$

(3) *If $\alpha \rho_{sv} \xi - \sqrt{\alpha(\alpha-1)} \xi < k < \alpha \rho_{sv} \xi + \sqrt{\alpha(\alpha-1)} \xi$,*

$$T^* = \frac{2}{\hat{v}(\alpha)} \left[\frac{\pi}{2} - \arctan \left(\frac{\alpha \rho_{sv} \xi - k}{\hat{v}(\alpha)} \right) \right], \quad (3.18)$$

where

$$\hat{v}(\alpha) = \sqrt{\alpha(\alpha-1) \xi^2 - (\alpha \rho_{sv} \xi - k)^2}.$$

$$(4) \text{ If } k \geq \alpha \rho_{sv} \xi + \sqrt{\alpha(\alpha-1)} \xi, \quad T^* = \infty. \quad (3.19)$$

PROOF See Appendix D (available online). \square

When the domestic and the foreign interest rates are constant, Proposition 3.4 examines moment boundedness in the Heston model. It is an extension of Proposition 3.1 in Andersen and Piterbarg (2007) from bounds on moments of order $\alpha > 1$ to bounds on all moments of order $\omega \in [\alpha, \alpha_1)$, for some $\alpha_1 > \alpha \geq 1$. We use this result to prove the strong convergence of the discretized discounted spot FX rate process.

PROPOSITION 3.5 *Let $\alpha \geq 1$. If $T < T^*$, there exist $\alpha_2 > \alpha$ and $\eta_\omega > 0$ such that, for all $\omega \in [1, \alpha_2)$,*

$$\sup_{\delta t \in (0, \eta_\omega)} \sup_{t \in [0, T]} \mathbb{E}[\bar{R}_t^\omega] < \infty, \quad (3.20)$$

where T^* is as given below.

$$(1) \text{ If } k < \alpha \rho_{sv} \xi + \frac{1}{2} \sqrt{\alpha(\alpha-1)} \xi, \quad T^* = \frac{1}{\alpha \rho_{sv} \xi + \sqrt{\alpha(\alpha-1)} \xi - k}. \quad (3.21)$$

$$(2) \text{ If } k \geq \alpha \rho_{sv} \xi + \frac{1}{2} \sqrt{\alpha(\alpha-1)} \xi, \quad T^* = \begin{cases} \infty & \text{if } \alpha = 1, \\ \frac{4(k - \alpha \rho_{sv} \xi)}{\alpha(\alpha-1)\xi^2} & \text{if } \alpha > 1. \end{cases} \quad (3.22)$$

PROOF See Appendix E (available online). \square

To the best of our knowledge, the boundedness of moments of discretization schemes for the Heston model and extensions thereof had not been established until recently (Cozma and Reisinger 2015a) – this gap in the literature was also identified in Kloeden and Neuenkirch (2012) – and our Proposition 3.5 is only the second result to address this issue. For the Heston model, Proposition 3.5 can be seen as an improvement of Proposition 3.9 in Cozma and Reisinger (2015a) due to the sharper conditions on the critical time.

3.2 The four-dimensional system

The strong mean square convergence of the discretized variance and domestic interest rate processes was established in Cozma and Reisinger (2015a, Proposition 3.5). First, we prove an equivalent result for the foreign interest rate.

PROPOSITION 3.6 *If $2k_f\theta_f > \xi_f^2$, then the process \bar{r}^f converges strongly in L^2 , ie,*

$$\lim_{\delta t \rightarrow 0} \sup_{t \in [0, T]} \mathbb{E}[|r_t^f - \bar{r}_t^f|^2] = 0. \quad (3.23)$$

PROOF See Appendix F (available online). \square

As will become clear from the proof, the Feller condition $2k_f\theta_f > \xi_f^2$, which ensures that the process r^f does not hit zero, allows us to control the potential growth of the absolute difference between the original and the discretized processes that comes from the sublinear correction term in the drift.

Second, we consider the logarithm of the process from (2.11) and examine its convergence properties. The formulas for the log-process, x , and its approximation, X , are

$$x_t = x_0 + \int_0^t (r_u^d - r_u^f - \frac{1}{2}v_u) du + \int_0^t \sqrt{v_u} dW_u^s, \quad (3.24)$$

$$\begin{aligned} X_t = x_0 + \int_0^t (\bar{r}_u^d - \bar{r}_u^f - \frac{1}{2}\bar{V}_u) du \\ + a_{11} \int_0^t \sqrt{\bar{V}_u} dW_u^1 + \sum_{j=2}^4 a_{1j} \int_0^t \sqrt{\bar{V}_u} \frac{\delta W_u^j}{\delta t} du. \end{aligned} \quad (3.25)$$

PROPOSITION 3.7 *If $2k_f\theta_f > \xi_f^2$, then the log-process converges uniformly in L^2 , ie,*

$$\lim_{\delta t \rightarrow 0} \mathbb{E} \left[\sup_{t \in [0, T]} |x_t - X_t|^2 \right] = 0. \quad (3.26)$$

PROOF See Appendix G (available online). \square

Third, we prove convergence of the discretized spot FX rate process.

PROPOSITION 3.8 *If $2k_f\theta_f > \xi_f^2$, then the process \bar{S} converges uniformly in probability, ie,*

$$\lim_{\delta t \rightarrow 0} \mathbb{P} \left(\sup_{t \in [0, T]} |S_t - \bar{S}_t| > \epsilon \right) = 0, \quad \forall \epsilon > 0. \quad (3.27)$$

PROOF See Appendix H (available online). \square

THEOREM 3.9 *Let $\alpha \geq 1$ and assume that the following conditions are satisfied:*

$$\left. \begin{aligned} k &> \alpha \rho_{sv} \xi + \max \{ \sqrt{\alpha(\alpha-1)} \xi, \frac{1}{4} \alpha(\alpha-1) T \xi^2 \}, \\ \frac{k_d^2}{2\xi_d^2} &> \frac{\alpha q_1(\alpha)}{q_1(\alpha)-1}, \quad \frac{2k_d}{T\xi_d^2} > \frac{\alpha q_2(\alpha)}{q_2(\alpha)-1} \quad \text{and} \quad 2k_f\theta_f > \xi_f^2. \end{aligned} \right\} \quad (3.28)$$

TABLE 1 The calibrated Heston parameters.

Parameter	Jessen and Poulsen	Elices and Giménez	Aït-Sahalia and Kimmel	Hurn <i>et al</i>
k	2.2200	1.1000	5.1300	1.9775
θ	0.0120	0.0097	0.0436	0.0376
ξ	0.1830	0.1400	0.5200	0.4568
ρ_{sv}	0.0634	0.1400	-0.7540	-0.7591

Column 2: for USD/EUR market data, January 2, 2004 to September 27, 2005 (Jessen and Poulsen 2013). Column 3: for EUR/USD market data of August 22, 2006 (Elices and Giménez 2013). Column 4: for the S&P 500 index between January 2, 1990 and September 30, 2003, using VIX data (Aït-Sahalia and Kimmel 2007). Column 5: for the S&P 500 index between January 2, 1990 and December 30, 2011, using two out-of-the-money options written on the index (Hurn *et al* 2014).

TABLE 2 The calibrated Cox–Ingersoll–Ross parameters.

Parameter	Driffill <i>et al</i>	Erismann	Brigo and Mercurio	Lafférs	Amin
$k_{d,f}$	0.0684	0.1104	0.3945	0.2820	0.1990
$\theta_{d,f}$	0.0161	0.0509	0.2713	0.0411	0.0497
$\xi_{d,f}$	0.0177	0.0498	0.0545	0.0058	0.0354

Column 2: to the three-month US Treasury bill yield between January 1964 and December 1998 (Driffill *et al* 2003). Column 3: to the US Treasury bill yield between October 1982 and April 2011 (Erismann 2011). Column 4: to the euro at-the-money caps volatility curve on January 17, 2000 (Brigo and Mercurio 2006). Column 5: to the euro overnight index average (Eonia) between January 1, 2008 and October 6, 2008 (Lafférs 2009). Column 6: using historical data for euro between January 1, 2001 and September 1, 2011 (Amin 2012).

Then the process converges strongly in L^α in the sense that

$$\lim_{\delta t \rightarrow 0} \sup_{t \in [0, T]} \mathbb{E}[|S_t - \bar{S}_t|^\alpha] = 0. \quad (3.29)$$

PROOF See Appendix I (available online). \square

Since the payoff of a typical FX or equity contract grows at most linearly in the exchange rate or the asset price, we only need to know the strong convergence in L^1 of the discounted process to deduce the convergence of the time-discretization error to zero. The following theorem can be generalized to the L^α case relatively easily, for all $\alpha \geq 1$, upon noticing that the critical time T^* from (3.16)–(3.19) is always greater than that from (3.21)–(3.22).

THEOREM 3.10 *If $2k_f\theta_f > \xi_f^2$ and $T < T^*$, then the discounted process converges strongly in L^1 , ie,*

$$\lim_{\delta t \rightarrow 0} \sup_{t \in [0, T]} \mathbb{E}[|R_t - \bar{R}_t|] = 0, \quad (3.30)$$

where

$$T^* = \begin{cases} \frac{1}{\rho_{sv}\xi - k} & \text{if } k < \rho_{sv}\xi, \\ \infty & \text{if } k \geq \rho_{sv}\xi. \end{cases} \quad (3.31)$$

PROOF The convergence in probability of the discounted process is a consequence of Proposition 3.8, by taking the domestic interest rate to be zero. The rest of the proof follows the argument of Theorem 3.9 closely and makes use of Propositions 3.4 and 3.5. \square

We can extend the convergence analysis from the four-dimensional Heston–CIR model to multi-factor short rates with CIR dynamics and a term structure, in which case Propositions 3.4–3.8, and hence Theorem 3.10, still hold, albeit with slightly modified proofs.

The condition $k \geq \rho_{sv}\xi$, also known as the “good correlation regime” (Jacquier and Martini 2011), is almost always satisfied in both FX and equity markets. This is because the speed of mean reversion k is usually larger than the volatility-of-volatility ξ , a fact clearly illustrated in Table 1. And even if this were not the case, a negative correlation between the underlying process and the variance, as is typically the case in equity markets (the so-called leverage effect), or a small absolute value of this correlation, as is typically the case in FX markets, would ensure the validity of the condition. Furthermore, the Feller condition $2k_f\theta_f > \xi_f^2$ in (3.28) for the foreign interest rate is generally satisfied in practice, a fact clearly illustrated in Table 2. We do not require a Feller condition for the stochastic volatility, and this is not always given in practice.

3.3 Option pricing

We conclude this section with a brief study on the convergence of mixed Monte Carlo/PDE estimators for computing FX option prices. Define the fair price of an option written on S ,

$$U = \mathbb{E} \left[\exp \left(- \int_0^T r_t^d dt \right) f(S) \right], \quad (3.32)$$

and its approximation under (2.11),

$$\bar{U} = \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}) \right], \quad (3.33)$$

where the payoff function f may depend on the entire path of the process and all expectations in this section are under the domestic risk-neutral measure \mathbb{Q} . The following theorem is concerned with the convergence of the time-discretization error to zero under the four-factor FX model, and can be extended straightforwardly to multi-factor

CIR short rates. The proof employs Propositions 3.4, 3.5 and 3.8 and Theorem 3.10. However, we omit it here because of its similarity to the proof of Cozma and Reisinger (2015a, Theorem 2.1) once the aforementioned results are established.

THEOREM 3.11 *Suppose that $2k_f\theta_f > \xi_f^2$. Then the following two statements hold.*

- (i) *The approximations to the values of the European put, the up-and-out call and any barrier put option defined in (3.33) converge as $\delta t \rightarrow 0$.*
- (ii) *If $T < T^*$, with T^* from (3.31), the approximations to the values of the European call, Asian options, the down-and-in/out and the up-and-in call option defined in (3.33) converge as $\delta t \rightarrow 0$.*

For European contracts, we can evaluate the conditional option price, ie, the inner-most expectation in (2.13), analytically. Hence, consider M simulations of the discrete paths of the variance and the interest rates and, for $1 \leq j \leq M$, let ω_j denote the j th sample. Then

$$\begin{aligned} \frac{1}{M} \sum_{j=1}^M \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}, \omega = \omega_j \right] \\ = \frac{1}{M} \sum_{j=1}^M \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d(j) dt \right) f(\bar{S}_T(j)) \mid \mathcal{G}_T^{f,d,v} \right] \end{aligned} \quad (3.34)$$

is the mixed Monte Carlo/PDE estimator of the European option price at time $t = 0$. The global error can be split into two parts:

$$\begin{aligned} \text{error} &= \mathbb{E} \left[\exp \left(- \int_0^T r_t^d dt \right) f(S_T) \right] \\ &\quad - \frac{1}{M} \sum_{j=1}^M \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}, \omega = \omega_j \right] \\ &= \left(\mathbb{E} \left[\exp \left(- \int_0^T r_t^d dt \right) f(S_T) \right] - \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \right] \right) \\ &\quad + \left(\mathbb{E} \left[\mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v} \right] \right] \right. \\ &\quad \left. - \frac{1}{M} \sum_{j=1}^M \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}, \omega = \omega_j \right] \right). \end{aligned}$$

The first term is the time-discretization error and the second term is the statistical error. The convergence to zero of the former was derived in Theorem 3.11 for European

put and call options: a result that can easily be extended to other financial derivatives, such as binary options. The convergence to zero of the latter follows from the Central Limit Theorem (see Glasserman 2003) upon noticing the following upper bound on the variance:

$$\begin{aligned} \text{Var} \left(\mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v} \right] \right) \\ \leq \mathbb{E} \left[\mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v} \right]^2 \right] \\ \leq \mathbb{E} \left[\exp \left(-2 \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T)^2 \right]. \end{aligned}$$

Assuming f has at most polynomial growth, we can employ Proposition 3.3 to deduce the finiteness of the variance under some conditions on the model parameters. In particular, if f is Lipschitz for $\alpha = 2$, Proposition 3.5 gives sharper sufficient conditions.

For path-dependent contracts, closed-form solutions are rarely available and we instead rely on finite differences to compute the conditional option price. Conditioning on the j th realization of the variance and interest rates paths, let $u_j^{f,d,v}(t, \bar{S}_t)$ and $\bar{u}_j^{f,d,v}(t, \bar{S}_t; P, L)$ be the solutions to the conditional PDE and to the associated finite-difference scheme, when a uniform mesh with P time steps and L spatial steps, respectively, is employed. It can be shown that the conditional PDE has a unique solution (see Section 7.1.2 in Evans 1998), which is, in fact, the conditional option price (see Theorem 7.3.1 in Shreve 2004). Then

$$\frac{1}{M} \sum_{j=1}^M \bar{u}_j^{f,d,v}(0, S_0; P, L) \quad (3.35)$$

is the mixed Monte Carlo/PDE estimator of the option price at $t = 0$. The global error can be split into three parts:

$$\begin{aligned} \text{error} &= \mathbb{E} \left[\exp \left(- \int_0^T r_t^d dt \right) f(S) \right] - \frac{1}{M} \sum_{j=1}^M \bar{u}_j^{f,d,v}(0, S_0; P, L) \\ &= \left(\mathbb{E} \left[\exp \left(- \int_0^T r_t^d dt \right) f(S) \right] - \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}) \right] \right) \end{aligned}$$

$$\begin{aligned}
& + \left(\mathbb{E} \left[\mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}) \mid \mathcal{G}_T^{f,d,v} \right] \right] \right. \\
& \quad \left. - \frac{1}{M} \sum_{j=1}^M \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}) \mid \mathcal{G}_T^{f,d,v}, \omega = \omega_j \right] \right) \\
& + \frac{1}{M} \sum_{j=1}^M (u_j^{f,d,v}(0, S_0) - \bar{u}_j^{f,d,v}(0, S_0; P, L)).
\end{aligned}$$

The first term is the time-discretization error, the second term is the statistical error and the third term is the finite-difference (FD) discretization error. The convergence to zero of the first term was derived in Theorem 3.11 for Asian and barrier options, and the convergence to zero of the second term is a consequence of the Central Limit Theorem (see the above discussion). However, the convergence of the third term, ie, of the finite-difference scheme, depends on the contract and the particular scheme employed.

4 VARIANCE REDUCTION ANALYSIS

In this section, we carry out a variance reduction analysis for European option valuation with the mixed Monte Carlo/PDE method under the four-factor FX model. However, the theory extends naturally to general interest rate dynamics. We use standard Monte Carlo with the log-Euler discretization as the reference method and define \hat{X} and \hat{S} to be the time-continuous approximations of x , defined in (3.24), and S , defined in (2.1), respectively. Then

$$\hat{X}_t = \hat{X}_{t_n} + (\bar{r}_{t_n}^d - \bar{r}_{t_n}^f - \frac{1}{2} \bar{V}_{t_n})(t - t_n) + \sqrt{\bar{V}_{t_n}} \Delta W_t^s, \quad (4.1)$$

where $\Delta W_t^s = W_t^s - W_{t_n}^s$ whenever $t \in [t_n, t_{n+1})$. Integrating (4.1) leads to

$$\hat{S}_t = S_0 \exp \left\{ \int_0^t (\bar{r}_u^d - \bar{r}_u^f - \frac{1}{2} \bar{V}_u) du + \int_0^t \sqrt{\bar{V}_u} dW_u^s \right\}. \quad (4.2)$$

We prefer the log-Euler scheme to the standard Euler scheme because it preserves positivity. Moreover, if the processes v , r^d , r^f are constant, then the first scheme is exact. Recall that

$$\begin{aligned}
\bar{S}_t = S_0 \exp \left\{ \int_0^t (\bar{r}_u^d - \bar{r}_u^f - \frac{1}{2} \bar{V}_u) du \right. \\
\left. + a_{11} \int_0^t \sqrt{\bar{V}_u} dW_u^1 + \sum_{j=2}^4 a_{1j} \int_0^t \sqrt{\bar{V}_u} \frac{\delta W_u^j}{\delta t} du \right\}.
\end{aligned}$$

Since \bar{V} is piecewise constant, we deduce that $\hat{S}_{t_n} = \bar{S}_{t_n}$ for all $0 \leq n \leq N$. The quantity that we want to estimate is the fair price of a European option with payoff function f , ie,

$$\Theta = \mathbb{E} \left[\exp \left(- \int_0^T r_t^d dt \right) f(S_T) \right]. \quad (4.3)$$

Then the corresponding standard and mixed Monte Carlo estimators are

$$\Theta_{\text{stdMC}} = \frac{1}{M} \sum_{j=1}^M \exp \left(- \int_0^T \bar{r}_t^d(j) dt \right) f(\hat{S}_T(j)), \quad (4.4)$$

$$\Theta_{\text{mixMC}} = \frac{1}{M} \sum_{j=1}^M \mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d(j) dt \right) f(\bar{S}_T(j)) \mid \mathcal{G}_T^{f,d,v} \right]. \quad (4.5)$$

Define

$$\text{Var}_{\text{stdMC}} = \text{Var}(\Theta_{\text{stdMC}}) = \frac{1}{M} \text{Var} \left(\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\hat{S}_T) \right) \quad (4.6)$$

and

$$\text{Var}_{\text{mixMC}} = \text{Var}(\Theta_{\text{mixMC}}) = \frac{1}{M} \text{Var} \left(\mathbb{E} \left[\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v} \right] \right). \quad (4.7)$$

Let the variance reduction factor (as in Dang *et al* 2015) and the standard deviation ratio be

$$\Gamma_{\text{var}} = \frac{\text{Var}_{\text{stdMC}}}{\text{Var}_{\text{mixMC}}} \quad \text{and} \quad \Gamma_{\text{dev}} = \sqrt{\Gamma_{\text{var}}}. \quad (4.8)$$

For convenience, we also define the discount factors

$$D = \exp \left(- \int_0^T r_t^d dt \right) \quad \text{and} \quad \bar{D} = \exp \left(- \int_0^T \bar{r}_t^d dt \right).$$

REMARK 4.1 From the “tower property” of conditional expectations, since $\hat{S}_T = \bar{S}_T$,

$$\mathbb{E}[\Theta_{\text{stdMC}}] = \mathbb{E}[\bar{D} f(\hat{S}_T)] = \mathbb{E}[\bar{D} f(\bar{S}_T)] = \mathbb{E}[\mathbb{E}[\bar{D} f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}]] = \mathbb{E}[\Theta_{\text{mixMC}}].$$

Therefore, the standard and mixed Monte Carlo estimators have the same discretization bias, ie,

$$\text{Bias}(\Theta_{\text{stdMC}}) = \text{Bias}(\Theta_{\text{mixMC}}). \quad (4.9)$$

REMARK 4.2 From the law of total variance we know that

$$\text{Var}(\bar{D} f(\bar{S}_T)) = \text{Var}(\mathbb{E}[\bar{D} f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}]) + \mathbb{E}[\text{Var}(\bar{D} f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v})]. \quad (4.10)$$

However, since $\hat{S}_T = \bar{S}_T$ and the variance is nonnegative, we deduce from (4.6) and (4.7) that the variance of the standard Monte Carlo estimator is greater than or equal to the variance of the mixed estimator, ie,

$$\text{Var}_{\text{stdMC}} \geq \text{Var}_{\text{mixMC}}. \quad (4.11)$$

Assuming a nontrivial payoff function f , equality occurs in (4.11) if and only if the second expectation on the right-hand side of (4.10) is zero, ie, if and only if \bar{S}_T is $\mathcal{G}_T^{\text{f,d,v}}$ -measurable. Since $\Sigma = AA^T$, after some straightforward calculations we find an equivalent condition:

$$\begin{aligned} a_{11} = 0 &\Leftrightarrow 1 - \rho_{\text{vd}}^2 - \rho_{\text{vf}}^2 - \rho_{\text{df}}^2 + 2\rho_{\text{vd}}\rho_{\text{vf}}\rho_{\text{df}} \\ &= \rho_{\text{sv}}^2(1 - \rho_{\text{df}}^2) + \rho_{\text{sd}}^2(1 - \rho_{\text{vf}}^2) + \rho_{\text{sf}}^2(1 - \rho_{\text{vd}}^2) \\ &\quad + 2\rho_{\text{sv}}\rho_{\text{sd}}(\rho_{\text{vf}}\rho_{\text{df}} - \rho_{\text{vd}}) + 2\rho_{\text{sv}}\rho_{\text{sf}}(\rho_{\text{vd}}\rho_{\text{df}} - \rho_{\text{vf}}) \\ &\quad + 2\rho_{\text{sd}}\rho_{\text{sf}}(\rho_{\text{vd}}\rho_{\text{vf}} - \rho_{\text{df}}). \end{aligned}$$

In particular, if the variance and the two interest rates are pairwise independent, then

$$a_{11} = 0 \Leftrightarrow \rho_{\text{sv}}^2 + \rho_{\text{sd}}^2 + \rho_{\text{sf}}^2 = 1.$$

Therefore, apart from this case, combining conditioning with Monte Carlo always reduces the variance of estimates. This is, however, to be expected since we eliminate the additional noise that comes from simulating the Brownian motion W^1 .

REMARK 4.3 Note that for any $2 \leq j \leq 4$, using the Ito isometry, we have

$$\text{Var} \left(\int_0^T \sqrt{\bar{V}_t} dW_t^j \right) = \mathbb{E} \left[\int_0^T \bar{V}_t dt \right]. \quad (4.12)$$

Moreover, using Cauchy's and Hölder's inequalities, Fubini's theorem and Remark 3.2 and Propositions 3.4 and 3.5 in Cozma and Reisinger (2015a), we can easily prove that

$$\lim_{\delta t \rightarrow 0} \frac{\mathbb{E}[\int_0^T \bar{V}_t dt]}{\text{Var}(\int_0^T \bar{V}_t dt)} = \frac{\mathbb{E}[\int_0^T v_t dt]}{\text{Var}(\int_0^T v_t dt)}. \quad (4.13)$$

From Dufresne (2001), we can compute the first two moments of the integrated square-root process explicitly. Hence, we find

$$\mathbb{E} \left[\int_0^T v_t dt \right] = \theta T + \frac{v_0}{k} - \frac{\theta}{k} + e^{-kT} \left(\frac{\theta}{k} - \frac{v_0}{k} \right) \quad (4.14)$$

and

$$\begin{aligned} &\text{Var} \left(\int_0^T v_t dt \right) \\ &= \frac{\xi^2}{k^2} \left[\theta T + \frac{v_0}{k} - \frac{5\theta}{2k} + 2e^{-kT} \left(\frac{\theta}{k} + \theta T - v_0 T \right) + e^{-2kT} \left(\frac{\theta}{2k} - \frac{v_0}{k} \right) \right] \end{aligned}$$

$$\begin{aligned}
&< \frac{\xi^2}{k^2}(1 + e^{-kT})\mathbb{E}\left[\int_0^T v_t dt\right] + \frac{\xi^2}{k^2}\left[\frac{\theta}{2k}e^{-kT}(4 + 2kT - e^{-kT} - 3e^{kT})\right] \\
&< \frac{\xi^2}{k^2}(1 + e^{-kT})\mathbb{E}\left[\int_0^T v_t dt\right].
\end{aligned} \tag{4.15}$$

Combining (4.12)–(4.15), we deduce that, for sufficiently small values of δt ,

$$\frac{\text{Var}(\int_0^T \sqrt{\bar{V}_t} dW_t^j)}{\text{Var}(\int_0^T \bar{V}_t dt)} > \frac{k^2}{\xi^2(1 + e^{-kT})} > \frac{k^2}{2\xi^2}. \tag{4.16}$$

The data in Tables 1 and 2 suggests that the interest rates have little impact on the variance of the mixed Monte Carlo estimator, and also that $k \gg \xi$ in both FX and equity markets. Hence, the stochastic integral on the left-hand side of (4.16) – part of the dividend yield q defined in (2.12) – contributes to the overall variance of the mixed estimator considerably more than the squared volatility σ^2 defined in (2.12). Therefore, we expect the minimum variance of the mixed estimator to be attained when all but the first term in q vanish, ie, when $a_{11} = 1$ and $a_{12} = a_{13} = a_{14} = 0$.

In practice, the volatility of interest rates in the CIR model is very small, ie, $\xi_{d,f} \ll 1$, a fact which can be observed in Table 2. Moreover, the volatility-of-volatility in the Heston model – calibrated to FX market data – is also small, ie, $\xi \ll 1$, and is significantly smaller than the rate of mean reversion, ie, $\xi \ll k$, a fact clearly illustrated in Table 1. Hence, the drift term in the square-root model for the variance dominates the diffusion term. Therefore, in the subsequent analysis we assume “almost deterministic” dynamics of the variance and interest rates. Let \bar{v} , \bar{v}^d and \bar{v}^f be the FTE discretizations of v , r^d and r^f corresponding to $\xi = \xi_d = \xi_f = 0$, ie, when the volatility-of-volatility parameters are equal to zero. Then

$$\begin{aligned}
\int_0^T \bar{r}_t^d dt &\approx \int_0^T \bar{v}_t^d dt, & \int_0^T \bar{r}_t^f dt &\approx \int_0^T \bar{v}_t^f dt, \\
\int_0^T \bar{V}_t dt &\approx \int_0^T \bar{v}_t dt, & \int_0^T \sqrt{\bar{V}_t} dW_t^s &\approx \int_0^T \sqrt{\bar{v}_t} dW_t^s.
\end{aligned}$$

REMARK 4.4 Suppose that W^s is independent of the Brownian motions W^v , W^d and W^f , ie, that $a_{11} = 1$ and $a_{12} = a_{13} = a_{14} = 0$. Employing (4.7) as well as the above assumption on the dynamics of the variance and the domestic and foreign interest rates, we obtain

$$\begin{aligned}
\text{Var}_{\text{mixMC}} &= \frac{1}{M} \text{Var} \left(\exp \left(- \int_0^T \bar{r}_t^d dt \right) \mathbb{E}[f(\bar{S}_T) \mid \mathcal{G}_T^{f,d,v}] \right) \\
&\approx \frac{1}{M} \exp \left(-2 \int_0^T \bar{v}_t^d dt \right) \text{Var}(\mathbb{E}[f(\bar{S}_T)]) \\
&= 0.
\end{aligned}$$

On the other hand, since $\hat{S}_T = \bar{S}_T$, we know from (4.6) that

$$\begin{aligned}\text{Var}_{\text{stdMC}} &= \frac{1}{M} \text{Var} \left(\exp \left(- \int_0^T \bar{r}_t^d dt \right) f(\bar{S}_T) \right) \\ &\approx \frac{1}{M} \exp \left(-2 \int_0^T \bar{\gamma}_t^d dt \right) \text{Var}(f(\bar{S}_T)).\end{aligned}$$

Assuming a nontrivial payoff f as well as a nonzero squared volatility v , since the variance of the mixed estimator is close to zero, this results in a substantial variance reduction.

REMARK 4.5 Let f be the European call option payoff and suppose that the dynamics of the variance and interest rates are “almost deterministic” (ie, their own volatility parameters are small compared to their mean-reversion speed). Then we can approximate the discounted payoff as follows:

$$\begin{aligned}P &= \exp \left(- \int_0^T \bar{r}_t^d dt \right) \\ &\quad \times \left(S_0 \exp \left\{ \int_0^T (\bar{r}_t^d - \bar{r}_t^f - \tfrac{1}{2} \bar{V}_t) dt + \int_0^T \sqrt{\bar{V}_t} dW_t^s \right\} - K \right)^+ \\ &\approx \exp \left(- \int_0^T \bar{\gamma}_t^d dt \right) \\ &\quad \times \left(S_0 \exp \left\{ \int_0^T (\bar{\gamma}_t^d - \bar{\gamma}_t^f - \tfrac{1}{2} \bar{\gamma}_t) dt + \int_0^T \sqrt{\bar{\gamma}_t} dW_t^s \right\} - K \right)^+. \quad (4.17)\end{aligned}$$

For now, assume that $a_{11} > 0$. For convenience, define the quantities

$$\left. \begin{aligned} \varrho &= \sqrt{1 - a_{11}^2}, & \tilde{\sigma} &= \sqrt{\int_0^T \bar{\gamma}_t dt}, \\ \tilde{D} &= \exp \left(- \int_0^T \bar{\gamma}_t^d dt \right), & F &= S_0 \exp \left\{ \int_0^T (\bar{\gamma}_t^d - \bar{\gamma}_t^f) dt \right\}, \end{aligned} \right\} \quad (4.18)$$

as well as

$$a = \frac{\varrho}{\sqrt{1 - \varrho^2}} \quad \text{and} \quad b = \frac{\log(F/K) + (0.5 + \varrho^2)\tilde{\sigma}^2}{\sqrt{1 - \varrho^2}\tilde{\sigma}}. \quad (4.19)$$

We use (4.17), the conditional option price formula in (2.14), and differentiate the variance of the mixed Monte Carlo estimator with respect to the parameter ϱ defined in (4.18) to find, using lengthy integration by parts,

$$\frac{\partial}{\partial \varrho} \text{Var}_{\text{mixMC}} \approx \frac{2}{M} \tilde{D}^2 F^2 \varrho \tilde{\sigma}^2 e^{\varrho^2 \tilde{\sigma}^2} \mathbb{E}[\Phi(aZ + b)^2], \quad (4.20)$$

where Φ is the standard normal CDF and $Z \sim \mathcal{N}(0, 1)$. Note that if $a_{11} < 1$, ie, if $\varrho > 0$, the right-hand side of (4.20) is strictly positive. This implies that the variance of the mixed estimator decreases as a_{11} increases, and attains its minimum when $a_{11} = 1$. In fact, we know from Remark 4.4 that

$$\text{Var}_{\text{mixMC}}(a_{11} = 1) \approx 0. \quad (4.21)$$

Moreover,

$$\mathbb{E}[\Phi(aZ + b)^2] = \Phi\left(\frac{b}{\sqrt{1+a^2}}\right) - 2T\left(\frac{b}{\sqrt{1+a^2}}, \frac{1}{\sqrt{1+2a^2}}\right), \quad (4.22)$$

where Owen's T function (Owen 1980) is

$$T(\beta, \vartheta) = \phi(\beta) \int_0^{\vartheta} \frac{\phi(\beta x)}{1+x^2} dx \quad \text{with } T(\beta, 1) = \frac{1}{2}\Phi(\beta) - \frac{1}{2}\Phi(\beta)^2, \quad (4.23)$$

and ϕ is the standard normal PDF. Note that

$$\frac{b}{\sqrt{1+a^2}} = \underbrace{\frac{1}{\tilde{\sigma}} \log \frac{F}{K}}_{\equiv \beta_1} + \underbrace{\left(\frac{1}{2} + \varrho^2\right) \tilde{\sigma}}_{\equiv \beta_2(\varrho)} \quad \text{and} \quad \frac{1}{\sqrt{1+2a^2}} = \underbrace{\sqrt{\frac{1-\varrho^2}{1+\varrho^2}}}_{\equiv \vartheta(\varrho)}. \quad (4.24)$$

According to the data in Table 1, $\theta \ll 1$ and so $\bar{\gamma} \ll 1$. Hence, for low maturities T we have $\tilde{\sigma} \ll 1$. From Remark 4.2, using (4.21), we know that the variance of the standard Monte Carlo estimator is obtained by integrating (4.20) over $[0, 1]$. Therefore,

$$\Gamma_{\text{var}} \approx \frac{\int_0^1 v[\Phi(\beta_1 + \beta_2(v)) - 2T(\beta_1 + \beta_2(v), \vartheta(v))] dv}{\int_0^1 v[\Phi(\beta_1 + \beta_2(v)) - 2T(\beta_1 + \beta_2(v), \vartheta(v))] dv}. \quad (4.25)$$

From (4.23) and (4.24), we deduce that, for all $v \in [0, 1]$,

$$\Phi(\beta_1 + 0.5\tilde{\sigma})^2 \leq \Phi(\beta_1 + \beta_2(v)) - 2T(\beta_1 + \beta_2(v), \vartheta(v)) \leq \Phi(\beta_1 + 1.5\tilde{\sigma}). \quad (4.26)$$

Therefore,

$$\frac{\Phi(\beta_1 + 0.5\tilde{\sigma})^2}{\Phi(\beta_1 + 1.5\tilde{\sigma})} \frac{1}{\varrho^2} \leq \Gamma_{\text{var}} \leq \frac{\Phi(\beta_1 + 1.5\tilde{\sigma})}{\Phi(\beta_1 + 0.5\tilde{\sigma})^2} \frac{1}{\varrho^2}. \quad (4.27)$$

The inequalities in (4.27) are approximate in the sense that the variance reduction factor is bounded from above and below by approximated quantities. Assuming that $\tilde{\sigma} \ll 1$, (4.27) becomes

$$\frac{\Phi(\beta_1)}{\varrho^2} \leq \Gamma_{\text{var}} \leq \frac{1}{\Phi(\beta_1)\varrho^2}. \quad (4.28)$$

Moreover, if we also assume that the option is deep in-the-money, then β_1 is relatively large (eg, $\beta_1 \geq 1$) and so $\sqrt{\Phi(\beta_1)} \approx 1$. Hence, we conclude that

$$\Gamma_{\text{dev}} \approx (1 - a_{11}^2)^{-1/2}, \quad \forall a_{11} > 0. \quad (4.29)$$

However, we know from Remark 4.2 that the standard deviation ratio is 1 when $a_{11} = 0$. Hence, (4.27)–(4.29) hold for all values of a_{11} . Interestingly, the deep in-the-money case is also the one where the payoff can be treated as smooth, and this also explains Figure 4. Note that Remark 4.5 is consistent with Remarks 4.2–4.4.

5 NUMERICAL RESULTS

In this section, we test the efficiency of the mixed Monte Carlo/PDE method by comparison with alternative numerical schemes for two derivatives: a European call and an up-and-out put option. For the former, we use the analytical formula (2.14) to compute the conditional option price and benchmark against standard Monte Carlo with a log-Euler discretization and, under a simple correlation structure, against the semi-analytical formula of Ahlip and Rutkowski (2013). For the latter, we use either finite differences or the perturbative formula of Fatone *et al* (2007) for the conditional option price, and Monte Carlo – with or without Brownian bridge – as the reference method.

First, we demonstrate the convergence of the mixed Monte Carlo/PDE method and find empirical convergence rates. Then, we investigate the sensitivity of the variance reduction factor (4.8) to changes in the model parameters and link the numerical results to the analysis in Section 4. Our platform for the numerical implementation was MATLAB 2012a. The machine configuration on which all numerical tests were conducted is as follows: Intel Core i3 CPU, M370, 2.40 GHz, 8.00 GB memory (RAM), 64-bit operating system running Windows 7 Professional.

Throughout this section, we assign the following values to the underlying model parameters as a base case, and vary a selection individually or jointly:

$$\begin{aligned} v_0 &= 0.0275, & r_0^d &= 0.0524, & r_0^f &= 0.0291, \\ k &= 1.70, & k_d &= 0.20, & k_f &= 0.32, \\ \theta &= 0.0232, & \theta_d &= 0.0475, & \theta_f &= 0.0248 \\ \xi &= 0.1500, & \xi_d &= 0.0352, & \xi_f &= 0.0317, \\ \rho_{sv} &= -0.10, & \rho_{sd} &= -0.15, & \rho_{sf} &= -0.15, \\ \rho_{vd} &= 0.12, & \rho_{vf} &= 0.05 & \rho_{df} &= 0.25. \end{aligned}$$

These values are consistent with empirical observations in FX markets and are close to the calibrated values in Tables 1 and 2. Also, the values of the correlations ρ_{sd} , ρ_{sf} and ρ_{df} are borrowed from Piterbarg (2006).

TABLE 3 Simulation results for a European call option.

Time steps	Discretization bias	Simulations	SD (stdMC)	SD (mixMC)
1	0.37231	1 000	0.50348	0.09342
2	0.08039	4 000	0.24126	0.04194
4	0.01775	16 000	0.12042	0.02023
8	0.00444	64 000	0.06071	0.00994
16	0.00160	256 000	0.02932	0.00487
32	0.00073	1 024 000	0.01507	0.00262

5.1 European call option

Let the spot, the strike and the maturity be $S_0 = 105$, $K = 100$ and $T = 1.5$. As an aside, note that FX option quotes are in terms of volatilities for a fixed delta and a fixed time to expiry, and not in terms of strikes. However, the strike price corresponding to a quoted volatility can easily be recovered using a conversion formula (Ahlip and Rutkowski 2013). Recall that Θ from (4.3) is the true option price, whereas Θ_{stdMC} from (4.4) and Θ_{mixMC} from (4.5) are the standard and mixed Monte Carlo estimators, respectively. In Table 3, we report the common discretization bias, ie, the time-discretization error, which is the same for both estimators by Remark 4.1, and the two standard errors, ie, the two standard deviations of the sample means, which we denote by SD and estimate using 10 000 samples. However, according to Ahlip and Rutkowski (2013), a closed-form solution for the European option price is not available under a full correlation structure. Hence, we need to find an accurate reference estimate Θ^* in order to evaluate the bias. We employ the mixed algorithm with $M = 2 \times 10^9$ simulations and $N = 200$ time steps to find $\Theta^* = 12.11968$. Then, the bias for a specific number of time steps is estimated using the reference estimate Θ^* , whose accuracy is discussed below, and a sufficiently large number of simulations, ie, $M = 2 \times 10^9$.

The time needed to obtain a call price estimate with 64 000 simulations and eight time steps is 0.25 seconds for the standard Monte Carlo method and 0.22 seconds for the mixed Monte Carlo method. Hence, the computational cost is 12% lower with the latter. This, however, is to be expected since we only simulate three of the four underlying processes.

The data in Table 3 indicates that the bias is small even when only a few time steps are used. This phenomenon can be explained by the fact that the Feller condition is satisfied, ie, $2k_\theta > \xi_\theta^2$, $2k_d > \xi_d^2$ and $2k_f > \xi_f^2$, and hence the discretizations of the variance and interest rate processes rarely hit zero. The simulation results in Table 3 confirm the square-root convergence of the statistical error and the first-

order convergence of the discretization error. Then, using extrapolation, we obtain an approximate root mean square error (RMSE) of the reference estimate:

$$\begin{aligned}\text{Bias}(\Theta^*) &\approx 1.168 \times 10^{-4}, \\ \text{SD}(\Theta^*) &\approx 0.593 \times 10^{-4} \Rightarrow \text{RMSE}(\Theta^*) \approx 1.310 \times 10^{-4}.\end{aligned}$$

This is equivalent to an RMSE of about 0.001% of the actual option price, suggesting that the reference estimate $\Theta^* = 12.11968$ is accurate to three decimal places.

Furthermore, the standard deviation is reduced by 83%, ie, is approximately six times lower with the mixed method. However, the variance reduction is not consistent across the range of possible values of the model parameters, and is most sensitive to changes in the correlations between the exchange rate and the squared volatility or the interest rates, ie, ρ_{sv} , ρ_{sd} and ρ_{sf} , the speed of mean reversion k and the volatility-of-volatility ξ .

The values $\rho_{vd} = 0.12$, $\rho_{vf} = 0.05$, $\rho_{df} = 0.25$ and $\rho_{sv}, \rho_{sd}, \rho_{sf} \in [-0.5, 0.5]$ lead to a valid (ie, positive definite) correlation matrix. The data in Figure 1 suggests that the highest variance reduction is achieved when the absolute values of the correlations ρ_{sv} , ρ_{sd} and ρ_{sf} are small or, more precisely, when $\rho_{sv} \approx -0.05$ and $\rho_{sd} \approx \rho_{sf} \approx 0$, in which case the standard deviation is reduced by a factor of 20. Hence, the best performance of the mixed estimator coincides with the independence of W^s from the Brownian motions W^v , W^d and W^f , an observation which is consistent with Remarks 4.3 and 4.4.

Furthermore, we may infer from Figure 1 that, in addition to the lower computational cost, the mixed Monte Carlo/PDE method outperforms standard Monte Carlo in terms of accuracy for all $\rho_{sv}, \rho_{sd}, \rho_{sf} \in [-0.5, 0.5]$, and that the variance reduction factor decreases as the absolute values of the correlations between the exchange rate and the squared volatility or the interest rates increase, two facts that were recognized in Remark 4.2.

Assuming that $\rho_{sd} = \rho_{sf}$, since $\rho_{vd} \approx \rho_{vf} \approx 0$, we may approximate

$$a_{11} \approx \sqrt{1 - \rho_{sv}^2 - \frac{2}{1 + \rho_{df}} \rho_{sd}^2} = \sqrt{1 - \rho_{sv}^2 - \frac{8}{5} \rho_{sd}^2}. \quad (5.1)$$

However,

$$\Gamma_{\text{dev}} \approx (1 - a_{11}^2)^{-1/2} \approx (\rho_{sv}^2 + 1.6\rho_{sd}^2)^{-1/2} \quad (5.2)$$

from Remark 4.5. Combined with the above observation on the location of the maximum variance reduction factor, we infer that the set of points (ρ_{sv}, ρ_{sd}) corresponding to $\Gamma_{\text{dev}} = \mu$ for some $\mu \geq 1$ approximately takes the form of an ellipse with the equation

$$\mu^2(\rho_{sv} + 0.05)^2 + 1.6\mu^2\rho_{sd}^2 = 1. \quad (5.3)$$

FIGURE 1 The standard deviation ratio Γ_{dev} from (4.8) with 4000 simulations and ten time steps, plotted against the correlation coefficients ρ_{sv} , ρ_{sd} and ρ_{sf} when the last two are equal.




TABLE 4 The estimated standard deviation ratio for different correlations.

ρ_{sv}	ρ_{sd}	Γ_{dev}
-0.40	-0.05	3.288
-0.40	0.05	3.088
-0.30	-0.20	3.387
-0.30	0.20	2.910
0.20	-0.20	2.951
0.20	0.20	3.127
0.30	-0.05	2.952
0.30	0.05	3.055

This confirms that the isolines in Figure 1 are approximately elliptical, a fact also illustrated in Table 4, where the correlation pairs (ρ_{sv}, ρ_{sd}) are chosen so that $\mu = 2.812$. The estimated standard deviation ratio along the contour line is 3.148 ± 0.239 , which is close to the theoretical value μ . This attests to the accuracy of our approximations and numerical results.

Next, we assume a partial correlation structure between the Brownian drivers such that the squared volatility dynamics are independent of the domestic and foreign interest rate dynamics, ie, $\rho_{vd} = \rho_{vf} = 0$. Using Remark 4.5 and computing a_{11}

TABLE 5 The empirical and theoretical standard deviation ratios for different correlations.

ρ_{sd}	ρ_{df}^*	Γ_{dev}	μ
-0.25	0.80	4.501	3.714
-0.20	1.00	5.301	4.472
-0.15	0.75	5.305	4.472
-0.10	0.50	5.224	4.472
-0.05	0.25	5.073	4.472
0	0	5.010	4.472
0.05	-0.25	4.850	4.472
0.10	-0.50	4.867	4.472
0.15	-0.75	4.554	4.472
0.20	-1.00	4.467	4.472
0.25	-0.80	3.788	3.714

explicitly from (2.2), we obtain

$$\Gamma_{dev} \approx \left[\rho_{sv}^2 + \frac{1}{1 - \rho_{df}^2} (\rho_{sd}^2 + \rho_{sf}^2 - 2\rho_{sd}\rho_{sf}\rho_{df}) \right]^{-1/2}. \quad (5.4)$$

Before, we fixed ρ_{df} and analyzed the variance reduction with respect to ρ_{sv} , ρ_{sd} and ρ_{sf} , when the latter two were equal. Now, we fix ρ_{sv} instead and focus on the effect of varying ρ_{df} on Γ_{dev} . To this end, we can easily show that

$$\frac{\rho_{sd}^2 + \rho_{sf}^2 - 2\rho_{sd}\rho_{sf}\rho_{df}}{1 - \rho_{df}^2} \geq \max\{\rho_{sd}^2, \rho_{sf}^2\}. \quad (5.5)$$


Assuming that ρ_{sd} and ρ_{sf} are not simultaneously zero, equality in (5.5) holds when

$$\rho_{df} = \rho_{df}^* \equiv \frac{\rho_{sf}}{\rho_{sd}} \mathbf{1}_{|\rho_{sd}| \geq |\rho_{sf}|} + \frac{\rho_{sd}}{\rho_{sf}} \mathbf{1}_{|\rho_{sd}| < |\rho_{sf}|}. \quad (5.6)$$

Upon its substitution into (5.4), we find a theoretical standard deviation ratio μ . Table 5 shows estimates of Γ_{dev} and compares them with μ when $\rho_{sv} = -0.10$, $\rho_{sf} = -0.20$, $\rho_{sd} \in [-0.25, 0.25]$ and $\rho_{df} \in [-0.85, 0.85]$. Note that these correlation values guarantee a symmetric positive definite correlation matrix. Hence, when ρ_{df}^* lies outside the interval of permitted values, we estimate Γ_{dev} using $\rho_{df} = \pm 0.85$ instead. The data in Table 5 suggests that the approximation (5.4) to the standard deviation ratio is fairly accurate, especially for a strongly negative correlation between the two interest rates.

By close inspection of the data in Figure 2, we infer that, for each $\rho_{sd} \in [-0.25, 0.25]$, the highest variance reduction is achieved for ρ_{df}^* as defined in (5.6). Hence, we conclude that μ exhibits the qualitative behavior of Γ_{dev} , so (5.4) provides

FIGURE 2 The standard deviation ratio with 1000 simulations and four time steps, plotted against the correlation coefficients ρ_{sd} and ρ_{df} when $\rho_{sv} = -0.10$, $\rho_{sf} = -0.20$ and $\rho_{vd} = \rho_{vf} = 0$.




a good approximation to the standard deviation ratio. In fact, our observations suggest that μ acts as a lower bound. We can extend these results to a full correlation structure between the Brownian drivers as long as ρ_{vd} and ρ_{vf} are close to zero, as seen before.

Suppose that the exchange rate dynamics are independent of the interest rate dynamics, ie, that $\rho_{sd} = \rho_{sf} = 0$, and define the optimal correlation ρ_{sv}^* to be the value corresponding to the maximum standard deviation ratio for a given volatility-of-volatility ξ . Then the data in Figure 3 suggests that the highest variance reduction is achieved when the absolute value of the correlation is small. In fact, we notice two things. First, the optimal correlation approaches zero as the volatility-of-volatility decreases, ie, $\lim_{\xi \rightarrow 0} \rho_{sv}^* = 0$, which is to be expected from Remark 4.3. Second, the standard deviation ratio at ρ_{sv}^* increases as the volatility-of-volatility decreases. In practice, $\xi_{d,f} \ll 1$ (see Table 2), so the two interest rates have little impact on the variance of the mixed Monte Carlo/PDE estimator. Therefore, since ρ_{sv}^* is close to zero, the variance comes mainly from σ defined in (2.12). On the other hand, smaller values of ξ result in a smaller variance of σ , and hence of the mixed estimator as well, which leads to a higher standard deviation ratio.

We observe a similar behavior when increasing the speed of mean reversion k or

FIGURE 3 The standard deviation ratio with 4000 simulations and ten time steps, plotted against the correlation ρ_{sv} and the volatility-of-volatility ξ when $\rho_{sd} = \rho_{sf} = 0$.




the long-run variance θ . Larger values of k result in a smaller variance of σ because of the mean-reverting property of the squared volatility, which ensures that the process returns to the long-run average quickly. On the other hand, larger values of θ produce higher volatilities, which then lead to an increase in the variance of the standard Monte Carlo estimator due to the larger diffusion term in the SDE driving the exchange rate process. We conclude the analysis by noting that values of ρ_{sd} and ρ_{sf} close to zero produce similar results, to some extent. However, when the absolute values of the two correlations are not small, the impact of σ , and hence of k , ξ and θ , on the variance of the mixed estimator is reduced.

The maximum in Figure 3 is attained around $\rho_{sv} = 0$ and $\xi = 0.05$, where the standard deviation ratio is $\Gamma_{dev} = 23$. Therefore, the same level of accuracy with the standard Monte Carlo method requires 529 times more simulations.

Next, we examine the variance reduction for different spots and maturities. The data in Figure 4 suggests that, unless the option is far out-of-the-money and the maturity is small, varying S_0 and T has little impact on the standard deviation ratio, which is approximately 2. Computing the option price with standard Monte Carlo means integrating the payoff function, which is not differentiable at the strike, whereas with the mixed Monte Carlo/PDE method we integrate the smooth conditional price. As the probability of a positive payoff decreases with S_0 and T , estimating it accurately with the former requires more simulations. Hence, the relative standard error of the standard Monte Carlo method increases as we go farther out-of-the-money or approach maturity, and the benefit of employing the mixed algorithm becomes clear.

FIGURE 4 The standard deviation ratio with 1000 simulations and four time steps, plotted against the spot S_0 and the maturity T when $\rho_{sv} = -0.60$.



For example, for a three-month call with a spot at 75% of the strike, we observe a variance reduction factor of 5275.

Furthermore, Figure 4 suggests that the approximation for the standard deviation ratio derived in (4.29) does not hold when the option is far out-of-the-money. For instance, we compute $a_{11} = 0.7893$, which gives a theoretical standard deviation ratio $\mu = 1.629$. When $S_0 = 105$ and $T = 0.5$, this is close to the estimated value, $\Gamma_{\text{dev}} = 1.832$. However, when $S_0 = 85$ and $T = 0.1$, we observe a much higher standard deviation ratio: $\Gamma_{\text{dev}} = 19.863$.

Suppose that the domestic and the foreign short rate dynamics are independent of each other, and also independent of the dynamics of the exchange rate and its volatility, and let $\rho_{sv} = -0.10$, as before. Moreover, suppose that the other model parameters, as well as the spot, the strike and the maturity, take the values listed at the beginning of this section. We use the mixed Monte Carlo/PDE method with $M = 8 \times 10^7$ simulations and $N = 200$ time steps, for an RMSE of about 0.001%, to obtain a call price estimate $\Theta' = 12.13621$, which is about 0.14% higher than the estimate corresponding to a full correlation of the factors, ie, $\Theta^* = 12.11968$. On the one hand, the postulated independence of the factors is critical from the point of view of analytical tractability (Ahlip and Rutkowski 2013), but can result in fairly different option prices. On the other hand, a full correlation structure leads to a richer model and a better fit to the observed market data. Finally, we test the accuracy of the mixed method and employ the semi-analytical pricing formula of Ahlip and Rutkowski (2013) to find the true option price, $\Theta = 12.13603$, and thus a relative

error of Θ' of about 0.0015%, which confirms that mixed Monte Carlo/PDE estimates are correct.

5.2 Up-and-out put option

Let the spot, the strike, the barrier and the maturity be $S_0 = 100$, $K = 105$, $B = 110$ and $T = 0.25$, and consider a continuously monitored up-and-out put option. We will first value the contract using the mixed Monte Carlo/PDE method. Hence, for a specific realization of the variance and interest rates paths, we compute the conditional option price, ie,

$$u(t, x) = \mathbb{E} \left[\exp \left(- \int_t^T \bar{r}_u^d du \right) (K - \bar{S}_T)^+ \mathbf{1}_{\max_{t \leq u \leq T} \bar{S}_u < B} \mid \mathcal{G}_T^{f,d,v}, \bar{S}_t = x \right]. \quad (5.7)$$


We know that u satisfies the following initial-boundary-value problem:

$$\left. \begin{aligned} \partial_t u + \mu_t x \partial_x u + \frac{1}{2} a_{11}^2 \bar{V}_t x^2 \partial_{xx} u - \bar{r}_t^d u &= 0, & \forall 0 < x < B, \quad t < T, \\ u(t, B) &= 0, & \forall t \leq T, \\ u(T, x) &= (K - x)^+, & \forall 0 \leq x < B. \end{aligned} \right\} \quad (5.8)$$

We solve the PDE backward in time from the initial condition, on a rectangular domain with $t \in [0, T]$ and $x \in [0.7S_0, B]$ that is discretized on a uniform grid with $N + 1$ temporal nodes and $L + 1$ spatial nodes. We selected this particular lower boundary of the spatial computational domain to reduce the number of spatial nodes, while making sure that the truncation error arising from our choice of the domain is negligible. We employ a central difference scheme to approximate the spatial derivatives and a linearity boundary condition (Tavella and Randall 2000) stating that $\partial_{xx} u = 0$ at the lower boundary, where the option is deep-in-the-money and the price can be regarded as linear in x . For convenience, we employ the same time grid as used for the discretization of the squared volatility and the interest rates. The final estimate of the barrier option price is a Monte Carlo average over a sufficiently large number of discrete trajectories of the Brownian motions W^2 , W^3 and W^4 .

Alternatively, we can use the perturbative formula of Fatone *et al* (2007) to approximate the conditional option price, and then a simple Monte Carlo average to estimate the outer expectation. Hence, we call this numerical scheme the mixed Monte Carlo/Pert method. Fatone *et al* (2007) approximate the up-and-out put option price in the Black–Scholes model with time-dependent parameters via a series expansion and provide explicit formulas for the first three terms, which involve some elementary and nonelementary transcendental functions. However, we will focus only on the zeroth-order approximation, because using a first-order correction term results in a hundredfold increase in computation time and hence in a poor performance of the

FIGURE 5 The log-plot of the time-discretization errors of the mixed Monte Carlo/PDE, the mixed Monte Carlo/Pert and the standard Monte Carlo methods against the time step.




scheme, compared with the mixed Monte Carlo/PDE method. Using $M = 4 \times 10^7$ simulations and $N = 200$ time steps to minimize the sampling and the discretization errors, respectively, we obtain a zeroth-order approximation: $\Theta^0 = 5.7700$.

Since a closed-form solution to the option pricing problem is not available, we need to find an accurate reference estimate Θ^* for the true option price Θ in order to compute the different errors of the numerical methods. Therefore, we use the mixed Monte Carlo/PDE algorithm with the Crank–Nicolson scheme, with $M = 4 \times 10^7$ simulations, $N = 200$ time steps and $L = 20$ spatial steps, to find $\Theta^* = 5.7631$. Hence, the approximation error of the mixed Monte Carlo/Pert method is $\Theta^0 - \Theta^* = 0.0069$, ie, about 0.1%.

In Figure 5, we report the time-discretization errors computed using the reference estimate Θ^* – whose accuracy is discussed below – or Θ^0 , and a sufficiently large number of simulations and spatial steps, ie, $M = 4 \times 10^7$ and $L = 20$. For the standard Monte Carlo and the mixed Monte Carlo/Pert methods, the time-discretization error is defined as the bias, whereas for the mixed Monte Carlo/PDE algorithm, due to our choice of the finite-difference grid, it includes the finite-difference (FD) time-discretization error. Henceforth, the term “discretization error” stands for the time-discretization error.

As an aside, note that crossings of the barrier are monitored only at discrete times by standard Monte Carlo. This gives rise to a monitoring error, which is included in the discretization error. Moreover, due to the knockout feature of the option, the true

FIGURE 6 The log-plot of the absolute space-discretization error of the mixed Monte Carlo/PDE method, against the spatial step.



price is smaller than the Monte Carlo estimate, which explains the strong positive bias displayed in Figure 5.

The data in Figure 5 suggests a square-root convergence of standard Monte Carlo and a first-order convergence of the mixed algorithms. On the other hand, we can use a Brownian bridge technique (see Glasserman 2003) to improve the first method and recover the first-order convergence. Indeed, the (red) Monte Carlo curve in Figure 5 would almost coincide with the (green) mixed Monte Carlo/Pert curve with the Brownian bridge correction. For instance, we calculated the discretization bias with $N = 8$ time steps to be 0.0075 for both Monte Carlo with Brownian bridge and mixed Monte Carlo/Pert.

In Figure 6, we report the space-discretization error computed using Θ^* and a sufficiently large number of simulations and time steps, ie, $M = 4 \times 10^7$ and $N = 200$. The data suggests a second-order convergence as well as a local minimum discretization error when the strike price lies halfway between two adjacent nodes, a technique known as grid-shifting (Tavella and Randall 2000). Hence, considering the first-order convergence of the time-discretization error (T-Err) and the second-order convergence of the finite-difference space-discretization error (S-Err) with the mixed Monte Carlo/PDE method, and using extrapolation, we obtain an approximate root mean square

error (RMSE) of the reference estimate:

$$\begin{aligned} \text{T-Err} &\approx 5.76 \times 10^{-4}, \\ \text{S-Err} &\approx 10.80 \times 10^{-4}, \\ \text{SD} &\approx 2.09 \times 10^{-4} \Rightarrow \text{RMSE} \approx 1.67 \times 10^{-3}. \end{aligned}$$

This is equivalent to an RMSE of about 0.03% of the actual option price, suggesting that the reference estimate $\Theta^* = 5.7631$ is accurate to two decimal places. Next, we compare the three numerical methods in terms of computation time for a given level of accuracy, in particular when the RMSE is at most 0.30% of the option price. First, using the empirical convergence rates determined above and extrapolation, we need $M = 2.5 \times 10^5$ simulations and $N = 800$ time steps, and hence a CPU time of 61.2 seconds, with the standard Monte Carlo method. Second, we reach this level of accuracy with the mixed Monte Carlo/PDE method when $M = 12\,000$, $N = 10$ and $L = 12$ in 2 seconds. Third, we need to employ the mixed Monte Carlo/Pert method with a zeroth-order approximation, with $M = 12\,000$ and $N = 10$, which takes 3.1 seconds.

Therefore, when the up-and-out put option price estimates need not be too accurate, eg, when one decimal place of accuracy is sufficient, the two mixed algorithms are comparable in terms of CPU time and efficiency and are considerably faster than standard Monte Carlo. However, a higher accuracy would require at least a first-order correction term in the mixed Monte Carlo/Pert approximation, making it highly time-consuming. We thus conclude that the mixed Monte Carlo/PDE method is the best of the three schemes.

We mentioned above that the Monte Carlo with Brownian bridge recovers the observed first-order convergence of the discretization error and the level of the bias from the mixed Monte Carlo/Pert method. The time-discretization error with the mixed Monte Carlo/PDE method is approximately 1.6 times larger, and includes the FD time-discretization error.

For a two-decimal-place accuracy, we fix one-hundred time steps and twenty spatial steps, such that the space- and time-discretization errors are about 0.02%. The time required to obtain a barrier option price estimate with 40 000 simulations is then 26 seconds for the mixed Monte Carlo/PDE method and 2.1 seconds for Monte Carlo with Brownian bridge (1.4 seconds for standard Monte Carlo). Therefore, the computational cost is 92% lower with the latter. In conclusion, due to the square-root convergence of the standard deviation, Γ_{dev} needs to be above 4.5 in order for the mixed Monte Carlo/PDE method to outperform Monte Carlo with Brownian bridge. Just as in the European call option case, the variance reduction is most sensitive to changes in the correlations between the exchange rate and the squared volatility or the interest rates.

FIGURE 7 The standard deviation ratio with 400 simulations, ten time steps and four spatial steps, plotted against the correlation coefficients ρ_{sv} , ρ_{sd} and ρ_{sf} when the latter two are equal.




Figure 7 exhibits similar characteristics to those of Figure 1. In particular, the highest variance reduction is achieved when $\rho_{sv} \approx 0.05$ and $\rho_{sd} \approx \rho_{sf} \approx 0$, in which case $\Gamma_{\text{dev}} = 40$. Based on a previous observation, this is equivalent to an eighty times lower computational effort with the mixed algorithm. A careful inspection of the data in Figure 7 suggests that the set of points (ρ_{sv}, ρ_{sd}) corresponding to $\Gamma_{\text{dev}} > 4.5$ can be described approximately by the following inequality:

$$(\rho_{sv} - 0.05)^2 + 1.6\rho_{sd}^2 < 4.2^{-2}, \quad (5.9)$$

ie, the inside of an ellipse, this being the set of parameters where the benefit of variance reduction outweighs the additional complexity of solving the conditional PDE numerically in this instance.

The variance reduction achieved by the mixed method results in computational savings (in the number of samples) by a factor hardly influenced by the desired accuracy. Conversely, higher accuracy of the finite-difference method can only be achieved by a larger number of mesh points. Hence, it would appear that for high enough accuracy the mixed method can never win over the standard Monte Carlo method. An asymptotic complexity gain of the mixed method for small errors would require that the PDE can be solved with constant effort independent of the desired accuracy (and for this effort to be outweighed by the reduced number of samples required for a given statistical error). Multilevel Monte Carlo methods (Giles and Reisinger 2012) recently developed for stochastic PDEs are designed precisely to achieve this goal,

Change OK?

by concentrating the dominant number of samples on the coarsest meshes while computing corrections on finer meshes with a vanishing number of paths. The application and numerical analysis in the present context is the subject of further research.

6 CONCLUSIONS

The numerical experiments carried out in Section 5 suggest that the mixed method outperforms both standard Monte Carlo and finite-difference methods under certain circumstances depending on the contract and the model parameters. When a closed-form solution for the conditional option price is available, we usually see a considerable improvement in both accuracy and computation time. When it is not, the mixed algorithm provides better accuracy at the expense of added computation time, and the set of parameter values where it outperforms the classical schemes is limited.

The analysis carried out in this paper is not restricted to the four-dimensional Heston–CIR model, but can be extended to higher-dimensional problems. For instance, multi-factor short rates may be considered as in Dang *et al* (2015), with CIR dynamics and a term structure, in which case the convergence and variance reduction analysis applies with some slight modifications of the proofs. Stochastic volatility accounts for volatility clustering, dependence in the increments and long-term smiles and skews, but gives rise to unrealistic short-term patterns in the implied volatility. Hence, in order to improve the behavior of the implied volatility for short maturities, the original model could be extended to a stochastic–local volatility model as in Cozma and Reisinger (2015a), which can easily be implemented for barrier option pricing with no extra computational effort, or an independent jump component could be added to the spot FX rate, in which case analytical formulas may be available for the conditional prices of European options. For instance, this is the case when the distribution of the jump size is normal (Merton 1976) or double-exponential (Kou 2002).

However, several open problems remain, such as the strong convergence rate of the discretization scheme or a finite-difference scheme with an observed second-order convergence in time for pricing the barrier option. In addition, examining the hedging parameters is also relevant, and we intend to pursue all these topics in our future research.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper. The research of Andrei Cozma is funded by the EPSRC and the support is gratefully acknowledged.

REFERENCES

- Ahlip, R., and Rutkowski, M. (2013). Pricing of foreign exchange options under the Heston stochastic volatility model and CIR interest rates. *Quantitative Finance* **13**(6), 955–966 (<http://doi.org/bmh2>).
- Aït-Sahalia, Y., and Kimmel, R. (2007). Maximum likelihood estimation of stochastic volatility models. *Journal of Financial Economics* **83**, 413–452 (<http://doi.org/b3c4dg>).
- Amin, H. H. N. (2012). Calibration of different interest rate models for a good fit of yield curves. Master's Thesis, Delft University of Technology.
- Andersen, L., and Piterbarg, V. (2007). Moment explosions in stochastic volatility models. *Finance and Stochastics* **11**(1), 29–50 (<http://doi.org/cbwdfw>).
- Ang, X. X. (2013). A mixed PDE/Monte Carlo approach as an efficient way to price under high-dimensional systems. Master's Thesis, University of Oxford.
- Brigo, D., and Mercurio, F. (2006). *Interest Rate Models: Theory and Practice*. Springer.
- Clark, I. J. (2011). *Foreign Exchange Option Pricing: A Practitioner's Guide*. Wiley.
- Cox, J., Ingersoll, J., and Ross, S. (1985). A theory of the term structure of interest rates. *Econometrica* **53**(2), 385–407 (<http://doi.org/cbb2pm>).
- Cozma, A., and Reisinger, C. (2015a). Convergence of an Euler discretisation scheme for the Heston stochastic–local volatility model with CIR interest rates. Working Paper, arXiv:1501.06084v3 [q-fin.CP].
- Cozma, A., and Reisinger, C. (2015b). Exponential integrability properties of Euler discretization schemes for the Cox–Ingersoll–Ross process. Working Paper, arXiv:1601.00919 [q-fin.CP].
- Dang, D.-M., Jackson, K. R., and Mohammadi, M. (2015). Dimension and variance reduction for Monte Carlo methods for high-dimensional models in finance. *Applied Mathematical Finance* **22**(6), 522–552 (<http://doi.org/bmh3>).
- Dereich, S., Neuenkirch, A., and Szpruch, L. (2012). An Euler-type method for the strong approximation of the Cox–Ingersoll–Ross process. *Proceedings of the Royal Society of London A* **468**, 1105–1115 (<http://doi.org/fphjf5>).
- Driffill, J., Kenc, T., and Sola, M. (2003). *An Empirical Examination of Term Structure Models with Regime Shifts*. Computing in Economics and Finance, Volume 65. Society for Computational Economics (<http://doi.org/dbmd7k>).
- Dufresne, D. (2001). The integrated square-root process. Research Paper 90. Centre for Actuarial Studies, University of Melbourne.
- Elices, A., and Giménez, E. (2013). Applying hedging strategies to estimate model risk and provision calculation. *Quantitative Finance* **13**(7), 1015–1028 (<http://doi.org/bmh4>).
- Erismann, M. (2011). Analytical propositions to evaluate contingent convertible capital. Master's Thesis, University of St Gallen.
- Evans, L. C. (1998). *Partial Differential Equations*, Graduate Studies in Mathematics, Volume 19. American Mathematical Society, Providence, RI.
- Fatone, L., Recchioni, M. C., and Zirilli, F. (2007). A perturbative formula to price barrier options with time-dependent parameters in the Black and Scholes world. *The Journal of Risk* **10**(2), 131–146 (<http://doi.org/bmh5>).
- Fischer, M., and Nappo, G. (2009). On the moments of the modulus of continuity of Itô processes. *Stochastic Analysis and Applications* **28**(1), 103–122 (<http://doi.org/cm6q3j>).

- Giles, M. (2008). Multilevel Monte Carlo path simulation. *Operations Research* **56**(3), 607–617 (<http://doi.org/fmrpfp>).
- Giles, M., and Reisinger, C. (2012). Stochastic finite differences and multilevel Monte Carlo for a class of SPDEs in finance. *SIAM Journal on Financial Mathematics* **3**(1), 572–592 (<http://doi.org/bmh6>).
- Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Stochastic Modelling and Applied Probability, Volume 53. Springer (<http://doi.org/bf7w>).
- Grzelak, L. A., and Oosterlee, C. W. (2011). On the Heston model with stochastic interest rates. *SIAM Journal on Financial Mathematics* **2**, 255–286 (<http://doi.org/bv6rd3>).
- Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* **6**(2), 327–343 (<http://doi.org/fg525s>).
- Higham, D. J., and Mao, X. (2005). Convergence of Monte Carlo simulations involving the mean-reverting square root process. *The Journal of Computational Finance* **8**(3), 35–62 (<http://doi.org/bmh7>).
- Hull, J., and White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance* **42**(2), 281–300 (<http://doi.org/bmh8>).
- Hunter, C. (2005). Hybrid derivatives. In *The Euromoney Derivatives and Risk Management Handbook*. Euromoney Institutional Investor, London.
- Hurn, A. S., Lindsay, K. A., and McClelland, A. J. (2014). Estimating the parameters of stochastic volatility models using option price data. *Journal of Business and Economic Statistics* **33**(4), 579–594 (<http://doi.org/bhwf>).
- Hutzenthaler, M., and Jentzen, A. (2015). *Numerical Approximations of Stochastic Differential Equations with Non-Globally Lipschitz Continuous Coefficients*. Memoirs of the American Mathematical Society, Volume 236. American Mathematical Society, Providence, RI (<http://doi.org/bmh9>).
- Jacquier, A., and Martini, C. (2011). Heston 2010. SSRN Working Paper (<http://doi.org/fztnkn>).
- Jessen, C., and Poulsen, R. (2013). Empirical performance of models for barrier option valuation. *Quantitative Finance* **13**(1), 1–11 (<http://doi.org/bmjbb>).
- Kloeden, P., and Neuenkirch, A. (2012). Convergence of numerical methods for stochastic differential equations in mathematical finance. In *Recent Developments in Computational Finance: Foundations, Algorithms and Applications*, Gerstner, T., and Kloeden, P. (eds). World Scientific.
- Kou, S. G. (2002). A jump–diffusion model for option pricing. *Management Science* **48**(8), 1086–1101 (<http://doi.org/fg58sf>).
- Lafférs, L. (2009). Empirical likelihood estimation of interest rate diffusion model. Master's Thesis, Comenius University, Bratislava.
- Lipp, T., Loeper, G., and Pironneau, O. (2013). Mixing Monte-Carlo and partial differential equations for pricing options. *Chinese Annals of Mathematics B* **34**(2), 255–276 (<http://doi.org/bmjcb>).
- Loeper, G., and Pironneau, O. (2009). A mixed PDE/Monte-Carlo method for stochastic volatility models. *Comptes Rendus Mathématique* **347**(9), 559–563 (<http://doi.org/b3bmnc>).

- Lord, R., Koekkoek, R., and van Dijk, D. (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance* **10**(2), 177–194 (<http://doi.org/bjnn3n>).
- McGhee, W. A. (2014). Pricing path dependent contracts in the presence of stochastic volatility: combining numerical integration, finite difference and conditional Monte Carlo. SSRN Working Paper (<http://doi.org/bhwg>).
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* **3**(1), 125–144 (<http://doi.org/d659w5>).
- Neuenkirch, A., and Szpruch, L. (2014). First order strong approximations of scalar SDEs defined in a domain. *Numerische Mathematik* **128**(1), 103–136 (<http://doi.org/bmjnd>).
- Owen, D. B. (1980). A table of normal integrals. *Communications in Statistics: Simulation and Computation* **B9**(4), 389–419 (<http://doi.org/dqzc7n>).
- Piterbarg, V. (2006). Smiling hybrids. *Risk* **19**(5), 66–71.
- Schöbel, R., and Zhu, J. (1999). Stochastic volatility with an Ornstein–Uhlenbeck process: an extension. *European Finance Review* **3**(1), 23–46 (<http://doi.org/bhxmdb>).
- Shreve, S. E. (2004). *Stochastic Calculus for Finance 2: Continuous-Time Models*. Springer Finance (<http://doi.org/bmjf>).
- Tavella, D., and Randall, C. (2000). *Pricing Financial Instruments: The Finite Difference Method*. Wiley.
- Van Haastrecht, A., and Pelsser, A. (2011). Generic pricing of FX, inflation and stock options under stochastic interest rates and stochastic volatility. *Quantitative Finance* **11**(5), 665–691 (<http://doi.org/bhc2b6>).
- Van Haastrecht, A., Lord, R., Pelsser, A., and Schrager, D. (2009). Pricing long-maturity equity and FX derivatives with stochastic interest rates and stochastic volatility. *Insurance: Mathematics and Economics* **45**(3), 436–448 (<http://doi.org/fjbt9z>).
- Willard, G. A. (1997). Calculating prices and sensitivities for path-independent derivative securities in multifactor models. *Journal of Derivatives* **5**(1), 45–61 (<http://doi.org/fg2kbn>).