

International Journal of Theoretical and Applied Finance
Vol. 17, No. 4 (2014) 1450024 (23 pages)
© World Scientific Publishing Company
DOI: 10.1142/S0219024914500241



EFFICIENT COMPUTATION OF EXPOSURE PROFILES FOR COUNTERPARTY CREDIT RISK

CORNELIS S. L. DE GRAAF

*Computational Science
University of Amsterdam, Science Park 904
Amsterdam, 1098 XH, The Netherlands
C.S.L.deGraaf@UvA.nl*

QIAN FENG

*CWI – The Center for Mathematics and Computer Science
Science Park 123, Amsterdam, 1098 XG, The Netherlands
Q.Feng@CWI.nl*

DRONA KANDHAI

*Quantitative Analytics, ING Bank
Bijlmerdreef 98, Amsterdam, 1102 CT, The Netherlands
Computational Science, University of Amsterdam
Science Park 904, Amsterdam, 1098 XH, The Netherlands
B.D.Kandhai@UvA.nl*

CORNELIS W. OOSTERLEE

*CWI – The Center for Mathematics and Computer Science
Science Park 123, Amsterdam, 1098 XG, The Netherlands
Delft University of Technology
DIAM – Delft Institute of Applied Mathematics
Mekelweg 4, Delft, 2628 CD, The Netherlands
C.W.Oosterlee@CWI.nl*

Received 24 January 2014

Accepted 11 March 2014

Published 13 June 2014

Three computational techniques for approximation of counterparty exposure for financial derivatives are presented. The exposure can be used to quantify so-called Credit Valuation Adjustment (CVA) and Potential Future Exposure (PFE), which are of utmost importance for modern risk management in the financial industry, especially since the recent credit crisis. The three techniques all involve a Monte Carlo path discretization and simulation of the underlying entities. Along the generated paths, the corresponding values and distributions are computed during the entire lifetime of the option. Option values are computed by either the finite difference method for the corresponding partial differential equations, or the simulation-based Stochastic Grid Bundling Method (SGBM), or by the COS method, based on Fourier-cosine expansions. In this research,

C. S. L. de Graaf et al.

numerical results are presented for early-exercise options. The underlying asset dynamics are given by either the Black–Scholes or the Heston stochastic volatility model.

Keywords: Expected exposure; potential future exposure; Bermudan options; Heston; numerical computation; finite differences; stochastic grid bundling method.

1. Introduction

During the recent financial crisis, several measures have been taken to make the financial industry more robust and resistant to financial shocks. The aim is to make individual participants better capable of handling possible credit events, like multiple simultaneous defaults. Regulations that are presented in the Basel III accords include a specific measure, named *Credit Valuation Adjustment* (CVA), which by definition is the difference between the risk-neutral valuation of a financial derivatives contract and the value which takes into account the possibility of a defaulting counterparty. In other words, CVA is the market value of *Counterparty Credit Risk* (CCR) (Zhu & Pykhtin, 2007).

The present research focuses on two important building blocks of CCR, *Expected Exposure* (EE) and *Potential Future Exposure* (PFE). The exposure is defined as the amount of money that may be lost if a counterparty defaults at a particular time, and cannot meet future payments that are agreed upon in the option contract. Because of the growing practical importance, there are recent articles from practitioners Antonov & Brecher (2012) and Antonov *et al.* (2011) that discuss the computation of the exposure. Since in particular the Heston model is widely used in practice, it is of interest to estimate the EE/PFE under the Heston model, and analyze the impact of a stochastic volatility model on the EE/PFE.

In over-the-counter (OTC) market, trades are settled directly between two parties and there is no third party to cover a possible huge loss because of a defaulting party. For European-style derivatives, without any opportunities to exercise contracts before the maturity date, it is obvious that CCR is important. When a counterparty defaults before the contract's maturity, the investment in the OTC option will be lost and the payoff will not be paid out. In the case of early-exercise options, such as Bermudan options, CCR is also relevant. In this case, although the holder has the right to exercise *at multiple moments during the life of the contract*, if an option is exercised because of financial distress of a counterparty, the return on investment is not as was originally expected, and therefore the option was most likely “mispriced”. Recent studies by Klein & Yang (2010) further elaborate on this issue.

In this paper, numerical methods are presented to keep track of the option values and their distributions during the life of the option contracts. All methods presented contain essentially two elements, a forward sweep for generating future scenarios and a backward sweep to calculate exposures along the generated asset paths. The forward Monte Carlo method generates the asset paths from initial time up to maturity. Along the paths, option values are determined at each exercise time.

Because of the complexity of this problem, efficient computation of the option prices is required. The COS Fourier option pricing method may seem a suitable candidate because of its speed and accuracy particularly for Lévy processes, see Fang (2010). Also the finite difference method, approximating solutions to partial differential equations, may be suitable as it typically results in approximate option prices for a grid of underlying values. This feature may be exploited in the EE context, as all grid points can then be used to generate option densities. The recent development of the Stochastic Grid Bundling Method (SGBM), which is a Monte Carlo-based method is particularly suitable for high-dimensional early-exercise options. Jain & Oosterlee (2012, 2013), is another candidate because it also rapidly converges and is accurate.

The setup of this paper is as follows. In Sec. 2, brief descriptions of CVA, exposure and other risk measures that are computed are presented. The underlying asset dynamics and the Monte Carlo discretization technique applied are also discussed. In Sec. 3, we describe the computation of the risk measures by the finite difference method, SGBM and the COS method, respectively. In Sec. 4, the methods are validated and compared and we present an assessment of the impact of stochastic volatility on the exposures. Finally, conclusions are presented in Sec. 5.

2. Problem Formulation

2.1. CVA and exposure of Bermudan options under Heston's model

CVA can be seen as the price of counterparty credit risk (Zhu & Pykhtin, 2007), while PFE is a measure for the potential loss. In other words, CVA depends on the Expected Exposure, while PFE is the loss given a fixed confidence interval. Both measures depend on the future distribution of exposure.

We will present methods for the computation of the exposure of Bermudan options under the Heston stochastic volatility asset dynamics, given by

$$\begin{aligned} dS_t &= rS_t dt + \sqrt{v_t} S_t dW_t^1, \\ dv_t &= \kappa(\eta - v_t) dt + \sigma \sqrt{v_t} dW_t^2, \\ dW_t^1 dW_t^2 &= \rho dt, \end{aligned} \quad (2.1)$$

where W_t^1 and W_t^2 are Wiener processes, correlated by parameter ρ , κ is the speed of mean reversion parameter in the CIR process for the variance, η represents the level of mean reversion, and σ is the so-called volatility of volatility parameter; r is the risk-free interest rate. The state of the process at time t_m is denoted by the pair (S_m, v_m) with S_m the price of the underlying and v_m the variance.

It is obvious that one is mainly interested in the case that a loss is positive (a negative loss may be a profit), and therefore the exposure at a future time $t < T$ is defined as

$$E(t) := \max(U(S_t, v_t, t), 0), \quad (2.2)$$

C. S. L. de Graaf et al.

where $U(S_t, v_t, t)$ is the (mark-to-market) value of a financial derivatives contract, like a call or put option.

The expected exposure (EE) at a future time t is given by

$$EE(t) := \mathbb{E}[E(t) | \mathcal{F}_0], \quad (2.3)$$

whereas $\theta = 97.5\%$ and $\theta = 2.5\%$ quantiles of the exposure distribution at time t are denoted by $PFE_{97.5\%}(t)$ and $PFE_{2.5\%}(t)$, respectively, are given as

$$PFE_{\theta}(t) = \inf\{x : \mathbb{P}(U(S_t, v_t, t) \leq x) \geq \theta\}. \quad (2.4)$$

Assuming independence^a between exposure and the counterparty's default probability, we can formulate the expression for CVA as follows (Gregory, 2010):

$$CVA(T) = (1 - \delta) \int_0^T D(t) EE(t) dPD(t), \quad (2.5)$$

where δ is the recovery rate, $D(t)$ is the risk-free discount factor, and $PD(t)$ denotes the default probability of the counterparty at time t (Gregory, 2010). Thus to summarize, the key elements here are: EE, recovery rate and the default probability. The EE is the amount one may lose at a specific time in the case of a defaulting counterparty, recovery rate is the percentage that can be recovered in case of default and the default probability is the probability that a counterparty may default in a certain period. All three elements are essential for the accurate computation of CVA. However, in this research, the focus is on the fast and accurate valuation of the exposure.

A Bermudan option is defined as an option where the buyer has the right to exercise at a set of (discretely spaced) time points. We denote the set of equally spaced exercise times by

$$\mathcal{T} = \{t_1, t_2, \dots, t_M\}, \quad (2.6)$$

where M denotes the number of exercise times, and the time difference is Δt . At the start of the option t_0 , exercise is not allowed.

At each exercise time, the exercise value, given by the payoff function, and the continuation value of the option are compared. The payoff function and the continuation value for the option at time t_m are, respectively, defined as:

$$\phi(S_m) = \max(\gamma(S_m - K), 0) \quad \text{with } \gamma = \begin{cases} 1 & \text{for a call,} \\ -1 & \text{for a put,} \end{cases} \quad (2.7)$$

$$c(S_m, v_m, t_m) = e^{-r\Delta t} \mathbb{E}[U(S_{m+1}, v_{m+1}, t_{m+1}) | (S_m, v_m)], \quad (2.8)$$

where $U(S_{m+1}, v_{m+1}, t_{m+1})$ is the option value at time t_{m+1} .

^aWrong way risk, which is an important notion in CVA depends particularly on nonzero correlation. This will be part of our future research.

Efficient Computation of Exposure Profiles for CCR

It is assumed that the holder of the option will exercise when the payoff value is higher than the continuation value, and then the contract terminates. At maturity t_M , the option value is equal to the payoff value.

The following recursive scheme can be setup to price a Bermudan option:

$$U(S_m, v_m, t_m) = \begin{cases} \phi(S_M) & \text{for } m = M, \\ \max[c(S_m, v_m, t_m), \phi(S_m)] & \text{for } m = 1, 2, \dots, M - 1, \\ c(S_0, v_0, t_0) & \text{for } m = 0. \end{cases} \quad (2.9)$$

By definition, the exposure of an option equals zero once the option is exercised; otherwise, the exposure is equal to the continuation value of the option. The Bermudan option exposure at a future time t_m can thus be formulated as:

$$E(t_m) = \begin{cases} 0 & \text{if exercised,} \\ c(S_m, v_m, t_m) & \text{if not exercised,} \end{cases} \quad m = 1, 2, \dots, M - 1. \quad (2.10)$$

In addition, $E(t_0) = c(S_0, v_0, t_0)$ and $E(t_M) = 0$.

The key point of calculating the exposure at time t_m is to determine the continuation value.

3. Numerical Methods to Compute Expected Exposure

In this section, three methods are presented to compute the expected exposure for Bermudan options under the Heston dynamics. All three methods can also be used to simply calculate the value of a Bermudan option at time t_0 . In combination with Monte Carlo forward path simulation, and based on the same common technique, they can be extended to value the exposure of Bermudan options.

3.1. General pricing approach

The market state depends on two random variables, (S_m, v_m) , at time point t_m and therefore the exposure $E(t_m)$ is also a stochastic variable. An option value distribution at future time points can be computed by generating scenarios, and therefore a Monte Carlo simulation is employed.

For the Monte Carlo simulation, the highly accurate Quadratic Exponential (QE) scheme (Andersen, 2008) is used here to generate the Heston stochastic volatility asset paths. Starting from simulated underlying values and variances, the exposures can be calculated by a backward valuation procedure. At each path, for each exercise time, the continuation value is calculated and compared to the exercise value on the path. When the exercise value is higher than the continuation value, the option is exercised at this path and the exposure for later time points is set to zero. At every time point, the resulting exposure values for all paths generate a distribution, as is illustrated in Fig. 1.

C. S. L. de Graaf et al.




Fig. 1. Monte Carlo paths simulated from left to right, an exposure distribution is obtained by calculating option values at the simulated states.

The essential technique of modeling the exposure of Bermudan options can be presented as follows:

- Generating scenarios/paths by Monte Carlo simulation;
- Calculate continuation/option values and the exercise values to decide whether to exercise or not;
- Set the exposure at each path as the continuation value if the option is not exercised; otherwise the exposure equals 0;
- Compute the empirical distribution of the exposure at each exercise time;
- Calculate EE, $PFE_{2.5\%}$ and $PFE_{97.5\%}$.

In the remainder of this section, we will describe three methods to calculate the required continuation/option values at the simulated paths.

3.2. The finite-difference-Monte-Carlo method

An often used option pricing technique is the finite difference method (Tavella & Randall, 2000). The method calculates option prices based on the option pricing partial differential equation, for an entire grid of underlying values and can therefore be easily used to compute the sensitivities (for example, the derivatives of the option prices w.r.t. the asset prices). In the scope of this research, the resulting grid of option values facilitates to determine distributions of option values at different time points. The method is called the *Finite-Difference-Monte-Carlo* (FDMC) method.

Efficient Computation of Exposure Profiles for CCR

Solving the Heston PDE to price European or American options is extensively studied, see, for example, Foulon & In't Hout (2010) and Haentjens & In't Hout (2013, 2013).

As described in Sec. 3.1, option values for all realized market states need to be computed. For European options it is well known that option value U satisfies the PDE:

$$\frac{\partial U}{\partial t} = \mathcal{A}U, \quad (3.1)$$

where in the Heston case the spatial differential operator \mathcal{A} is given by

$$\begin{aligned} \mathcal{A}U = & \frac{1}{2}\sigma^2v\frac{\partial^2U}{\partial v^2} + \rho\sigma vS\frac{\partial^2U}{\partial S\partial v} + \frac{1}{2}vS^2\frac{\partial^2U}{\partial S^2} \\ & + (\kappa[\eta - v])\frac{\partial U}{\partial v} + rS\frac{\partial U}{\partial S} + \frac{\partial U}{\partial t} - rU. \end{aligned} \quad (3.2)$$

For American options, a linear complementarity problem is solved. Using the payoff function (2.7), the option value in this case satisfies:

$$\frac{\partial U}{\partial t} \geq \mathcal{A}U, \quad (3.3a)$$

$$U(S) \geq \phi(S), \quad (3.3b)$$

$$(U - \phi(S)) \left(\frac{\partial U}{\partial t} - \mathcal{A}U \right) = 0, \quad (3.3c)$$

with one equality sign in either (3.3a) or (3.3b). Note that for each $t \in [t_0, T]$ the option can be exercised. A discrete version easily results in the pricing of a Bermudan-style option.

In this research, the Brennan–Schwartz (1977) algorithm is used, which is a well-known technique from literature. At each exercise time, this method first solves inequality (3.3a) as an equality, after which the option value is taken to be the maximum of this value and the exercise value.

The boundary conditions used are stated in Table 1. Note that at the $v = 0$ boundary a so-called degenerated boundary condition is imposed which is obtained by substituting $v = 0$ in (3.1).

The schemes used for discretizing (3.2) in asset and variance directions are second-order accurate central schemes or one-sided second-order schemes needed at boundaries.

Table 1. Heston model boundary conditions for a European put option.

Boundary	Value
$S \rightarrow \infty$	$U = 0$
$S = 0$	$U = K$
$v \rightarrow \infty$	$\frac{\partial^2 U}{\partial v^2} = 0$
$v = 0$	$\frac{\partial U}{\partial t} - rU + rS\frac{\partial U}{\partial S} + \kappa\eta\frac{\partial U}{\partial v} = 0$

C. S. L. de Graaf et al.

The option price is computed backwards in time, from maturity T back to time t_0 . The equations that need to be solved as a result of the finite difference discretization are linear systems of equations. Such a system of equations can be represented as a matrix-vector problem where the operators are represented by matrices and the (intermediate) solutions by vectors. For the time integration scheme, a particular Alternating Direction Implicit (ADI) scheme, namely the Hundsdorfer–Verwer scheme is employed, which exhibits second-order convergence in time. Next to that, due to the splitting of matrices, it involves the inversion of tridiagonal matrices while in general for (fully) implicit schemes, the matrices are not tridiagonal and may have several nonzero diagonals. For more details we refer to Foulon & In't Hout (2010).

To ensure that all paths are contained in the computational domain of the finite difference technique, the boundaries S_{\max} and v_{\max} are prescribed such that all Monte Carlo path values at all time points are contained.

The paths attain values that are, most likely, not grid points of the finite difference grid. From this grid, specific option values are determined by interpolation. Because this interpolation may introduce errors at every point, second-order accurate *spline* interpolation is used.

In general, only a small part of the discretized grid is a region of interest, therefore one can concentrate grid points in that region. This is done by stretching the grid so that a nonuniform grid results, applied in the variance as well as in the asset dimension (Haentjens & In't Hout, 2013). As here we need option values at each exercise time for many combinations of spot and variance values, nonuniformity is even more important when we compute exposure.

Because tests show that the impact of the spot dimension on the error is highest, the nonuniform grid in Haentjens & In't Hout (2013) is slightly adjusted. The grid employed is a combination of a uniform and a nonuniform grid. An interval $[S_{\text{left}}, S_{\text{right}}]$ containing K is introduced in which the mesh is uniform. We choose:

$$S_{\text{left}} = \lambda K \quad \text{and} \quad S_{\text{right}} = K, \quad (3.4)$$

where $\lambda \in [0.3, 0.7]$ can be chosen depending on the quantity that needs to be computed (PFE or EE). An accurate computation of EE requires accurate pricing around the mean, which implies a high value of λ , whereas for accurate computation of quantiles, an accurate computation of extreme values is needed for which a smaller value of λ should be chosen. So, when we compute the PFE, the dense region is shifted toward the outer regions of the domain.

Outside the interval $[S_{\text{left}}, S_{\text{right}}]$, the grid follows a hyperbolic sine function with:

$$\begin{aligned} \xi_{\min} &= \xi_0 = \sinh^{-1} \left(\frac{S_{\text{left}}}{d_1} \right), \\ \xi_{\text{int}} &= \frac{S_{\text{right}} - S_{\text{left}}}{d_1}, \\ \xi_{\max} &= \xi_{m1} = \xi_{\text{int}} + \sinh^{-1} \left(\frac{S_{\max} - S_{\text{right}}}{d_1} \right), \end{aligned}$$

Efficient Computation of Exposure Profiles for CCR

where d_1 is a scaling parameter, and $\xi_{\min} < 0 < \xi_{\text{int}} < \xi_{\max}$. Now, to construct a grid of $m_1 + 1$ points, the *nonuniform* adjustment $0 = s_0 < s_1 < \dots < s_{m_1} = S_{\max}$ can be constructed via a *uniform* grid of $m_1 + 1$ points between ξ_{\min} and ξ_{\max} : $\xi_{\min} = \xi_0 < \xi_1 < \dots < \xi_{m_1} = \xi_{\max}$ and the function g :

$$g(\xi_i) =: s_i \quad i = 0, \dots, m_1, \text{ where}$$

$$g(\xi_i) = \begin{cases} S_{\text{left}} + d_1 \sinh(\xi_i) & \text{if } \xi_{\min} \leq \xi_i < 0, \\ S_{\text{left}} + d_1 \xi_i & \text{if } 0 \leq \xi_i \leq \xi_{\text{int}}, \\ S_{\text{right}} + d_1 \sinh(\xi_i - \xi_{\text{int}}) & \text{if } \xi_{\text{int}} < \xi_i \leq \xi_{\max}. \end{cases}$$

Smaller values of d_1 result in a lower density in $[S_{\text{left}}, S_{\text{right}}]$, whereas higher values of d_1 will result in a higher density of grid points in this interval.

For the v direction, more points at the boundary $v = 0$ are desired and for larger values of v , the mesh can be less dense. Let m_2 be the number of points to be considered and d_2 another scaling parameter. Define equidistant points $\nu_0 < \nu_1 < \dots < \nu_{m_2}$, given by $\nu_j = j \cdot \Delta\nu$, with $\Delta\nu = \frac{1}{m_2} \sinh^{-1}(\frac{v_{\max}}{d_2})$, for $j = 0, \dots, m_2$. Now the grid $0 = v_0 < v_1 < \dots < v_{m_2} = v_{\max}$ is defined by: $v_j = d_2 \sinh(\nu_j)$, $j = 0, \dots, m_2$.

These grids are smooth in the sense that there are real-valued constants C_0, C_1 and C_2 such that:

$$C_0 \Delta\xi \leq \Delta s_i \leq C_1 \Delta\xi \quad \text{and} \quad |\Delta s_{i+1} - \Delta s_i| \leq C_2 (\Delta\xi)^2.$$

When the finite difference method is used to price a single option, only a single grid point at initial time is used. The FDMC method however uses a large portion of the grid points for option pricing at all exercise times which makes this method computationally attractive.

3.3. Stochastic-grid-bundling method

The *Stochastic-Grid-Bundling* Method (SGBM) is a Monte Carlo method that combines *regression*, *bundling* and *simulation*. It was proposed by Jain & Oosterlee (2012, 2013) for pricing multi-dimensional Bermudan options under Black-Scholes dynamics. The SGBM method generates a direct estimator, a lower bound for the option value, as well as an optimal early-exercise policy. Here, we extend the SGBM method from Jain & Oosterlee (2013) towards the Heston model, and exposure distributions along the time horizon are naturally obtained.

Suppose we deal with a Bermudan option with tenor T and M exercise dates. First a stochastic grid is generated, i.e. we generate H paths of the underlying under the Heston model. It is easy to see that the option value at time $t_M = T$ is equal to the corresponding payoff value, which gives us the initial setting for the SGBM method at each path.

C. S. L. de Graaf et al.

At time t_k , $k = M - 1, \dots, 1$, these paths are clustered into β bundles, based on their stock and variance values. The bundle set at time t_m is denoted by $\{\mathcal{B}^{p,m}\}_{p=1}^\beta$. Paths within the same bundle are assumed to share some common properties. We adapt the so-called multi-dimensional recursive bifurcation bundling method (Jain & Oosterlee, 2013) by means of a rotation.

Within each bundle, basis functions $\{g_k(S, v)\}_{k=0}^B$ are defined for the regression of option values.^b The essential idea in SGBM is that, for paths in the p th bundle $\mathcal{B}^{p,m}$ at time t_m , a set of coefficients $\{\alpha_k^{p,m}\}_{k=0}^B$ exists, so that for the option values of these paths at time t_{m+1} , the following relationship holds

$$U(S_{m+1}, v_{m+1}, t_{m+1}) \approx \sum_{k=0}^B \alpha_k^{p,m} g_k(S_{m+1}, v_{m+1}). \quad (3.5)$$

When the option values $U(S_{m+1}, v_{m+1}, t_{m+1})$ at the stochastic paths are determined, the coefficient set $\{\hat{\alpha}_k^{p,m}\}_{k=0}^B$ can be obtained by regression. Equation (3.5) can be substituted into (2.8) which gives us:

$$\begin{aligned} c(S_m, v_m, t_m) &= e^{-r\Delta t} \mathbb{E}[U(S_{m+1}, v_{m+1}, t_{m+1}) | (S_m, v_m)] \\ &\approx e^{-r\Delta t} \mathbb{E}\left[\sum_{k=0}^B \hat{\alpha}_k^{p,m} g_k(S_{m+1}, v_{m+1}) | (S_m, v_m)\right] \\ &= e^{-r\Delta t} \sum_{k=0}^B \hat{\alpha}_k^{p,m} \mathbb{E}[g_k(S_{m+1}, v_{m+1}) | (S_m, v_m)] \\ &= e^{-r\Delta t} \sum_{k=0}^B \hat{\alpha}_k^{p,m} f_k(S_m, v_m). \end{aligned} \quad (3.6)$$

When the functions $\{f_k(\cdot, \cdot)\}_{k=0}^B$ are known, the continuation values at time t_m can be computed, and, subsequently, the option values at time t_m can be obtained with the scheme in (2.9). At time t_0 , we deal with one bundle, as all paths originate from (S_0, v_0) , and the option value at time t_0 is equal to the continuation value $c(S_0, v_0, t_0)$. In this way, option values are calculated backward in time from t_M to t_0 . By (2.10), the exposure at each path along the time horizon is calculated as a by-product. We are thus able to determine the empirical exposure distribution at each time point for the calculation of EE and PFE.

We choose the basis functions $\{g_k(S, v)\}$ such that analytic formulas for their expectations are available. Analytic formulas of these expectations bring exact information into the recursive procedure. Under the Heston dynamics, the basis functions are chosen as

$$g_k(S, v) := (\log(S))^k, \quad (3.7)$$

where $k = 0, \dots, B$.

^bFor the definition of basis function see Longstaff & Schwartz (2001).

When $k = 0$, the basis function is the constant; when $k > 1$, the expectation of g_k is the k th moment of $\log(S)$.

There is a well-known relationship between the moments of the model and its characteristic function. The joint characteristic function of Heston's model is available (see, for example, Fang & Oosterlee, 2011) and thus analytic formulas for the expectation $f_k(\cdot, \cdot)$ can be derived. Note, however, that the stochastic grid is based on (S, v) -values; only the regression is based on log-variables.

The optimal exercise strategy is determined by comparing the immediate exercise value and the continuation value, and the exercise decision is made when the immediate exercise value is highest. We store the strategy and the corresponding realized cash flows at each path during the backward procedure. The option value can be calculated as the mean of the discounted cash flow as in Longstaff & Schwartz (2001).

The SGBM method has some advantages compared to the well-known Longstaff–Schwartz method, although both are based on regression. The Longstaff–Schwartz method uses only the “in-the-money” paths to get the optimal stopping strategy and the corresponding cash flow. In SGBM, all paths are used and the optimal stopping strategy is merely a by-product. By applying bundling, the approximation of the linear coefficients can be optimized locally. Furthermore, information from the model dynamics is included by application of the analytic formulas for the expectation of the basis functions, whereas the Longstaff–Schwartz method only employs the dynamics in the path generation.

3.4. The COS-Monte-Carlo method

In the third computational method, we combine the generated stochastic MC grid with the COS method, introduced in Fang & Oosterlee (2011). Based on the same stochastic grid as, for example, in Sec. 3.3, the COS method is used for the calculation of the continuation values at each path along the time horizon. We call this combined method the *COS-Monte-Carlo* (CMC) method.

As in Fang & Oosterlee (2011), we work in the log-domain, denoted by $(x, u) := (\log(S), \log(v))$. Suppose that the path values (x_m, u_m) at time t_m are known. We can write the joint density function at t_{m+1} , conditioned on values at t_m , as

$$f_{x,u}(x_{m+1}, u_{m+1} | x_m, u_m) = f_{x|u}(x_{m+1} | x_m, u_{m+1}, u_m) \cdot f_u(u_{m+1} | u_m), \quad (3.8)$$

where $f_{x|u}(\cdot)$ is the conditional log-stock density, and $f_u(\cdot)$ the conditional log-variance density. Notice that here we have $f_u(u_{m+1} | x_m, u_m) = f_u(u_{m+1} | u_m)$.

The continuation value defined in (2.8) at time t_m is the expectation of the option value at time t_{m+1} w.r.t the joint density function. One can choose a proper integration range $[a, b] \times [a_v, b_v]$ in log-stock domain and log-variance domain, so that the integral can accurately be approximated. We refer to Fang & Oosterlee (2011) for details on the definition of this range based on initial state (x_0, u_0) .

In this paper, we are not only concerned with the option value at time t_0 , but also with continuation values along the time axis. To assure accuracy, we need a

C. S. L. de Graaf et al.

common integration range which is sufficiently large for all paths at each exercise time. We define it as

$$\begin{aligned} [a, b] &:= \bigcup_{h,m=1}^{H,M} [a^{h,m}, b^{h,m}], \\ [a_v, b_v] &:= \bigcup_{h,m=1}^{H,M} [a_v^{h,m}, b_v^{h,m}], \end{aligned} \quad (3.9)$$

where $[a^{h,m}, b^{h,m}]$ and $[a_v^{h,m}, b_v^{h,m}]$ are the ranges for the log-stock and log-variance domains, respectively, for the h th path at t_m , by the suggestions in Fang (2010) and Fang & Oosterlee (2011).

The integration can now be written as follows:

$$\begin{aligned} c(x_m, u_m, t_m) &\approx e^{-r\Delta t} \int_{a_v}^{b_v} f_u(u_{m+1} | u_m) \cdot \left[\int_a^b U(x_{m+1}, u_{m+1}, t_{m+1}) f_{x|u} \right. \\ &\quad \left. \times (x_{m+1} | x_m, u_{m+1}, u_m) dx_{m+1} \right] du_{m+1}. \end{aligned} \quad (3.10)$$

An analytic formula for the log-variance density $f_u(\cdot | \cdot)$ is available (see Fang & Oosterlee, 2011), and the conditional log-stock density $f_{x|u}(\cdot | \cdot, \cdot, \cdot)$ can be recovered from the corresponding characteristic function (Fang & Oosterlee, 2011) by applying the COS expansion. The recovered density of the log-stock process from the characteristic function (Fang & Oosterlee, 2011) is given by

$$\begin{aligned} f_{x|u}(x_{m+1} | x_m, u_{m+1}, u_m) &\approx \frac{2}{b-a} \cdot \sum'_{n=0}^{N-1} \operatorname{Re} \left\{ \Phi \left(\frac{n\pi}{b-a}; u_{m+1}, u_m \right) e^{in\pi \frac{x_m - a}{b-a}} \right\} \\ &\quad \times \cos \left(n\pi \frac{x_{m+1} - a}{b-a} \right), \end{aligned} \quad (3.11)$$

where \sum' indicates that the first term is multiplied by $\frac{1}{2}$, $\operatorname{Re}(\cdot)$ returns the real part of the value, the function Φ is defined as $\Phi(\frac{n\pi}{b-a}; u_{m+1}, u_m) := \Phi(\frac{n\pi}{b-a}; 0, u_{m+1}, u_m)$, which is the characteristic function of the log-stock process.

We use the Gaussian-quadrature rule (Fang & Oosterlee, 2011) for the approximation of the outer integral in (3.10). The log-variance integral range is discretized on a grid, denoted by $\{\varsigma_j, \varsigma_0 = a_v, \varsigma_J = b_v\}_{j=0}^J$. The characteristic function of the log-stock process at time t_{m+1} conditioned on the log-variance and (x_m, u_m) is denoted by $\Phi(\omega; x_m, \varsigma_j, u_m)$.

By interchanging the inner integral and the summation obtained by the COS expansion, the continuation value can be written as

$$\begin{aligned} c(x_m, u_m, t_m) &\approx e^{-r\Delta t} \sum_{j=0}^{J-1} w_j f_u(\varsigma_j | u_m) \cdot \sum'_{n=0}^{N-1} A_{n,j}(t_{m+1}) \\ &\quad \times \operatorname{Re} \left\{ \Phi \left(\frac{n\pi}{b-a}; \varsigma_j, u_m \right) e^{in\pi \frac{x_m - a}{b-a}} \right\}, \end{aligned} \quad (3.12)$$

Efficient Computation of Exposure Profiles for CCR

where the w_j are the weights of the quadrature nodes ς_j , $j = 0, 1, \dots, J - 1$; $A_{n,j}(t_{m+1})$ is the Fourier-cosine coefficient defined as follows:

$$A_{n,j}(t_{m+1}) = \frac{2}{b-a} \left[\int_a^b U(x_{m+1}, \varsigma_j, t_{m+1}) \cos \left(n\pi \frac{x_{m+1} - a}{b-a} \right) dx_{m+1} \right]. \quad (3.13)$$

At time t_M , the values of the coefficients $A_{n,j}(t_M)$ can easily be obtained as the option value at time t_M equals the payoff value. The expression for $A_{n,j}(t_M)$ becomes

$$A_{n,j}(t_M) = \begin{cases} G_n(a, 0) & \text{for a put,} \\ G_n(0, b) & \text{for a call,} \end{cases} \quad (3.14)$$

where the G_n -functions are the cosine coefficients of the payoff function, given as:

$$G_n(l, u) = \frac{2}{b-a} \int_l^u \phi(y) \cos \left(n\pi \frac{y-a}{b-a} \right) dy. \quad (3.15)$$

Fortunately, for specific payoffs, an analytic formula for $G_n(l, u)$ is available (see Fang & Oosterlee, 2011).

From maturity t_M a backward recursive calculation can be used to obtain all coefficients $A_{n,j}(t_m)$.

For the computation of the EE, however, interpolation is needed for the calculation of the continuation value at each time step for each path. We use interpolation to reduce the calculation costs, because the calculation of function $\Phi(\frac{n\pi}{b-a}; \varsigma_j, u_m)$ takes significant CPU time. Instead of calculating the value of function u_m for each path and time step t_m , we compute an $N \times J \times J$ matrix Φ with elements $(\Phi(\frac{n\pi}{b-a}; \varsigma_j, \varsigma_p))$, $j, p = 0, \dots, J, n = 1, \dots, N$. For a fixed p , we extract the corresponding 2D slice from the 3D matrix to calculate the continuation value $c(x_m, \varsigma_p, t_m)$ at time t_m .

One can easily determine vector $\{c(x_m, \varsigma_p, t_m)\}_{p=0}^J$, storing the continuation values of each path on the variance grid by matrix calculation. As in the FDMC method, by spline interpolation, an accurate continuation value $c(x_m, u_m, t_m)$ at each path can be obtained.

After the calculation of the continuation value, the exposure can be determined easily by applying the formulas in Sec. 2.1.

Compared to the COS method for pricing option values, the CMC method is significantly slower when the number of MC paths is high. One reason is that, at each exercise time, an additional calculation of the continuation value is performed, for which interpolation is required for each path. Next to that, as we need to choose a wider integration range to assure accuracy for EE and PFE, we also need a large number of Fourier-cosine terms and variance grid points to get converged results, which has a significant impact on the computational speed.

At the same time, the CMC method maintains the very high accuracy of the COS method. The errors due to the truncated integration ranges, the quadrature and the propagation error have been discussed in Fang & Oosterlee (2011). The

C. S. L. de Graaf et al.

error of the spline interpolation on the variance grid is small when J is sufficiently large as the continuation value is a continuous function of the variance. Because of the high accuracy, we will use the results of the CMC method as reference values in the discussion of the numerical results.

4. Numerical Results

In this section, we start with an assessment of the impact of stochastic volatility on the EE and PFE profiles. Next, we consider a detailed analysis of the convergence and accuracy of the methods by means of numerical experiments.

As there are no exact values available for the exposure of Bermudan put options under the Heston stochastic volatility model, we will use the converged results of the COS method as reference values.^c As mentioned in Sec. 3.4, the COS method is a highly accurate method for pricing Bermudan options. When valuing the exposure, the high accuracy is maintained as long as the integration range is chosen properly (see Sec. 3.4). We reduce the impact of Monte Carlo noise in the comparative analysis by using 10^5 paths.

To investigate the proposed three methods, three different sets of parameters are tested, see Table 2. These test cases were used recently in Fang & Oosterlee (2011), Haentjens & In't Hout (2013), and Vellekoop & Nieuwenhuis (2009) and reference values are thus available for individual option prices. Moreover, in these test cases, we stress the parameters of the stochastic volatility process by considering different levels for the initial variance, the mean-reversion parameters, vol-of-vol and correlation parameters. These parameters are chosen such that in Tests A and C the well-known Feller condition is satisfied, while in Test B it is not.^d Apart from the different settings for the model parameters, we consider different maturities, interest rates and moneyness levels.

Table 2. Parameter sets for Tests A, B and C.

	Test A	Test B	Test C
Spot (S)	10	100	9
Strike (K)	10	100	10
Interest (r)	0.04	0.04	0.10
Exercise Times	50	50	50
Initial Vol ($\sqrt{v_0}$)	0.5745	0.1865	0.2500
Tenor (T)	0.25	0.25	1.00
Mean Reversion (κ)	0.80	1.15	5.00
Mean Var (η)	0.3300	0.0348	0.1600
Vol of Var (σ)	0.700	0.459	0.900
Correlation (ρ)	0.10	-0.64	0.10

^cThe convergence of the COS method has been discussed in Fang & Oosterlee (2011), and we will set the number of Fourier terms to $N = 2^9$ and the number of the variance grid points to be $J = 2^9$.

^dIt is known that when the Feller condition is not satisfied, the variance process can become zero and numerical methods can suffer from this issue.

4.1. Comparison of Black–Scholes to Heston to assess impact of stochastic volatility on exposure

Here it is shown that stochastic volatility clearly has an impact on exposure profiles. It is most significant for the $PFE_{97.5\%}$ quantile in the tests considered. We restrict the analysis to Tests A and B because in Test C the mean reversion level is not equal to the initial variance and thus it is not clear which level to use for the variance in the Black–Scholes model. In Fig. 2, the results are plotted for the parameters from Tests A and B.

In general, independent of the underlying dynamics, the plots show that the EE starts at the initial option value, after that, the level drops because of the early exercise possibility. The PFEs also start at the initial option value, because at this stage there is no uncertainty, i.e. the minimum value for which the probability is higher than a specific benchmark that equals this initial value. Starting from $t = 0$, $PFE_{2.5\%}$ drops to zero soon while $PFE_{97.5\%}$ is always higher than the EE. Due to the early exercise possibility, paths will “terminate” i.e. exercise will take place so that more than 2.5% of the values are equal to zero soon. With the same argument, the minimum value for which 97.5% of the prices are lower is much higher and only drops at a later stage as more and more paths are being exercised.

When the results for Black–Scholes are compared to Heston, one can conclude that the most significant difference is for $PFE_{97.5\%}$, in both cases. The difference in $PFE_{97.5\%}$ is a factor 10 times larger than the difference for EE and $PFE_{2.5\%}$. Intuitively this makes sense, due to the fact that the mean reversion level is equal to the constant variance level in the Black–Scholes model, the EE is not heavily affected. However, since the volatility is stochastic, extreme cases may occur more frequently (with the parameters chosen), resulting in fatter tails of the distribution that have a significant impact on PFE.




Fig. 2. EE and PFE profiles under the Black–Scholes and Heston model, differences are significant for 97.5% PFE.

C. S. L. de Graaf et al.

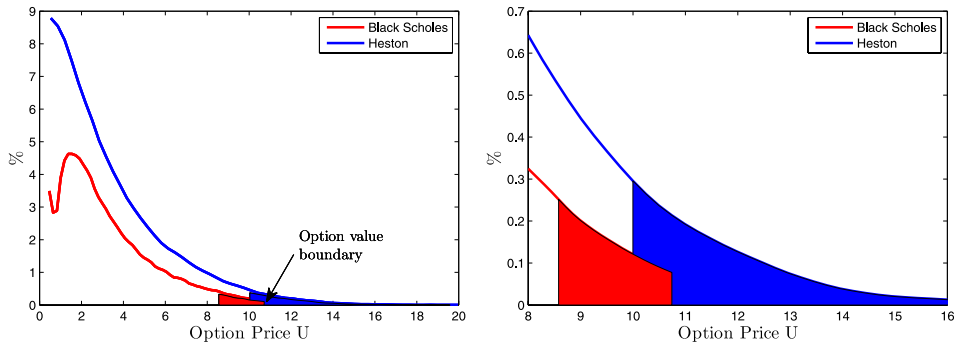


Fig. 3. Distribution of option values under Heston and Black–Scholes for Test B at exercise time 17 of 50. The $PFE_{97.5\%}$ is shaded in red (BS) or blue (Heston), respectively, in the right plot the axes are changed to make the boundary more clear.

The early exercise value also depends on the volatility so that for any path, there is a different exercise value. From Fig. 3, it can be seen that the distribution computed under the Black–Scholes dynamics is chopped off at a specific maximum option value, whereas the distribution under the Heston dynamics has a smoothly varying tail. The mass that is originally in the cut of tail in the Black–Scholes case is here located at the left-side boundary.

Although these results show that stochastic volatility has an impact on the exposure profiles, a more rigorous analysis, based on market calibrated parameters instead of model parameters, will be subject of further research. Here instead we will focus on the accuracy and numerical convergence of our proposed methods.

For the analysis of the accuracy and convergence of the three proposed methods, we concentrate on Tests B and C. The results for EE and $PFE_{97.5\%}$ obtained for the different methods are shown in Fig. 4 and Table 3 (the results for $PFE_{2.5\%}$ are not shown in the table). The methods are tested on a set of $H = 10^5$ generated Monte




Fig. 4. EE, $PFE_{2.5\%}$ and $PFE_{97.5\%}$, for 10^5 paths and 50 exercise times.

Efficient Computation of Exposure Profiles for CCR

Table 3. Relative L_2 difference between three methods for 10^5 paths.

Error	Measurement	Test B	Test C
EE	$\ FD - COS\ _2$	5.8864e-03	1.8966e-03
	$\ SGBM - COS\ _2$	8.0645e-03	1.4943e-03
PFE _{97.5%}	$\ FD - COS\ _2$	1.4133e-02	2.7113e-03
	$\ SGBM - COS\ _2$	4.5856e-03	3.9364e-03

Carlo paths. In the following subsections, the convergence and error behavior of the FDMC method and SGBM is discussed.

4.2. Error FDMC

The error for pricing options with the finite difference method for the Heston PDE is extensively studied, see, for example, Foulon & In't Hout (2010). The error is mainly introduced near the boundaries, but can be controlled by a combination of a large number of grid points and the use of a nonuniform grid. In all finite difference computations, the grids are nonuniform, as discussed in Sec. 3.2. The free parameter λ in (3.4) to determine the region $[\lambda K, K]$ in spot direction is determined depending on the quantity that is being measured: For PFE a smaller value is desired, whereas for EE the value is larger, in any case $\lambda \in [0.3, 0.7]$. The variance grid is very dense around the $v = 0$ boundary, independent of the measured quantity. The number of grid points in spot (S) and variance (v) directions are denoted by m_1 and m_2 , respectively. By experiment we know that the numerical error is dominated by the error in spot direction, and therefore the number of grid points in the S direction is chosen as $m_1 = 2m_2$. With this fixed ratio, the decay in error is measured by decreasing a generic measure Δs defined as $\Delta s := \frac{1}{m_2}$. If we decrease Δs by increasing the number of grid points, the numerical convergence is second order when we price a single Bermudan option in Test C, whereas it is almost second order for Test B, see Fig. 5. The grid is chosen to be very dense in the region of the initial market parameters (S_0, v_0) , and the price is extracted from the grid by accurate spline interpolation.

When EE and PFE _{θ} are computed, multiple prices at each exercise time are needed. In this case, interpolation is needed for each path and exercise time which is expected to have an impact on the error. To investigate the scale of this error, the same convergence tests are done as in the single option case. In this case, the finite difference solution is compared to the semi-analytic CMC method described in Sec. 3.4. The same random scenarios are used for computing the EE for the CMC and the FDMC methods. As shown in Fig. 6, in both Tests B and C, the convergence of the error is similar for EE as it is for pricing a single option. The decrease of the error is of second order in the number of grid points in Test C and almost second order in Test B.

For PFE, in Test B the convergence from the start is only first order. In this test, the Feller condition not satisfied. The mesh used in this case has a dense region

C. S. L. de Graaf et al.




Fig. 5. Convergence plots by increasing the number of grid points in space for a single Bermudan option. The relative L_2 norm is used to measure the difference with the reference COS value.




Fig. 6. Convergence plots by increasing the number of grid points in space for computing EE and PFE. The relative L_2 norm is used to measure the difference with the reference value obtained by the CMC method.

around the strike, whereas for a PFE computation, the strike region is generally not of highest relevance. To enhance the accuracy of the PFE, the nonuniform grid in S -direction can be adjusted (which we leave for later study). Because PFE and EE are mostly computed independently, a conclusion is that measuring PFE or EE would imply using two different grids.

The convergence with respect to Δt is not presented in this research because tests show that the error is dominated by the spatial error.

4.3. Error SGBM

Here we focus on the convergence of SGBM regarding the option value, the EE and PFE. We use five basis functions (including the constant) defined in Eq. (3.7). In

Efficient Computation of Exposure Profiles for CCR

the recursive bifurcation bundling method, an essential property is that the number of bundles must be of the form 4^j , $j = 0, 1, 2, \dots$, for details we refer to Jain & Oosterlee (2013). The bundling scheme is slightly adapted to deal with the two-dimensional Heston dynamics.

In the tests, a large number of paths $H = 10^5$, and bundles $\beta = 4^4$ are chosen.

In Sec. 3.3, there are two ways of calculating the option value at time t_0 . One way is to estimate the coefficient set over all paths at time t_1 and to apply regression at time t_0 (the so-called direct estimator); the other is to store the optimal strategy and take the mean of the discounted cash flow (the so-called path estimator):

- The results calculated directly from the set of Monte Carlo paths is called *direct estimator* results;
- The results calculated by the second set of paths, but with the coefficients from the first set of paths is called the *path estimator*.

The numerical results for the path estimator should be similar to the results for the direct estimator. Table 4 presents the difference between the direct estimator and the path estimator for EE and PFE. Again the error is measured in the relative L_2 norm. We can see that the difference between the two results is only of order 10^{-3} .

The option value is the maximum value obtained among all possible stopping rules, indicating that the option value calculated by the “optimal” strategy will be less than or equal to the real option value. This provides a criterion for convergence. The result calculated by the optimal strategy will be the lower bound of the Bermudan option value.

We examine the convergence of the Bermudan option value w.r.t the number of bundles for SGBM. The tests are done for 10 simulations, and the results are presented in Fig. 7. We take the regression results of the direct estimator and the results of the optimal strategy of the path estimator for comparison. As we can see, in both Tests B and C, the results of the path and the direct estimator resemble each other better when the number of bundles increases. The two results are very close to the COS reference value for $\beta = 4^4$, see Table 5.

In addition to the convergence of the Bermudan option value, we examine the convergence of EE and PFE in Fig. 8. The results of the CMC method is used as the reference value. With the same set of 10^5 generated paths, we increase the number of bundles from 1 to 4^4 for the calculation of the SGBM method. It shows that the error

Table 4. The difference between the direct estimator and path estimator for EE and PFE. The number of bundles equals 4^4 and the number of paths is 10^5 .

	EE	PFE
Test B	2.8541e-03	6.3035e-03
Test C	1.9462e-03	5.9252e-03

C. S. L. de Graaf et al.

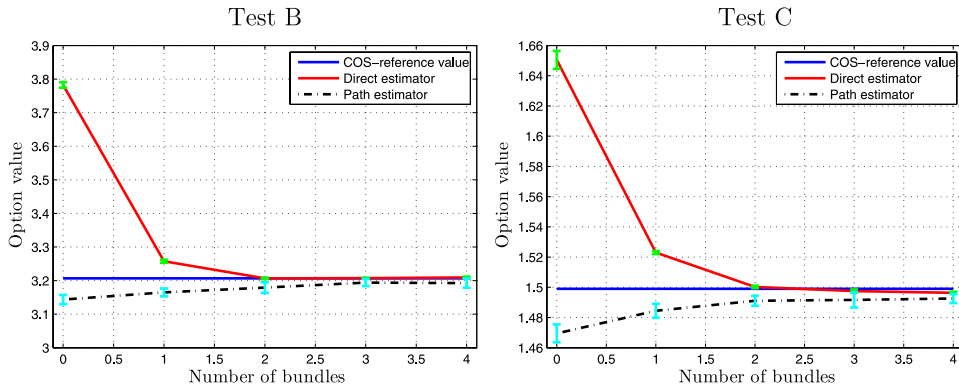


Fig. 7. Convergence with respect to the number of bundles $\beta = 4^j$, for the Bermudan option value; the total number of paths $H = 10^5$; the reference value in Test B equals 3.2066 and in Test C 1.4990. The red dashed line is the direct estimator, blue is the COS reference value, and the black dashed line the path estimator.

Table 5. The difference between the direct and the path estimator for a Bermudan option value when the number of bundles equals 4^4 . The results are computed via 10 simulations (s.e. is standard error).

	COS (reference)	Direct estimator (s.e.)	Path estimator (s.e.)
Test B	3.2066	3.2091 (2.8613e-03)	3.1924 (1.3768e-02)
Test C	1.4990	1.4964 (7.2086e-04)	1.4926 (3.0376e-03)




Fig. 8. Convergence with respect to the number of bundles 4^j ; the total number of paths equal 10^5 .

decreases when increasing the number of bundles. The EE results exhibit a higher accuracy than the $PFE_{97.5\%}$ results, when the number of bundles is equal to 4^4 .

The convergence of the EE and PFE, w.r.t the number of paths, is examined in Fig. 9. We choose the number of bundle equal to 4^3 , and increase the average

Efficient Computation of Exposure Profiles for CCR




Fig. 9. Convergence with respect to the average number of paths in each bundle; the number of bundles is equal to 4^3 . The black line indicates an asymptotic convergence with $1/\sqrt{\tilde{H}_\beta}$, where \tilde{H}_β denotes the average number of paths per bundle.

number of paths in each bundle. The differences of EE and PFE between direct estimator and path estimator are compared. The average number of paths in each bundle is increased from 25 to 2000. It shows that the difference between the path and the direct estimator decreases when the average number of paths in each bundle (i.e. the total number of paths) increases.

These results support the fact that SGBM converges (to the reference values) for Bermudan options, EE and $\text{PFE}_{97.5\%}$.

5. Conclusion

In this paper, three different approaches for computing exposure profiles within the context of counterparty credit risk are presented. The underlying asset exposure is driven by the Heston stochastic volatility model and Bermudan put options are priced. In all three methods, scenarios are generated by a Monte Carlo scheme and option values are priced at each path at each exercise time. The pricing procedure is done by either the developed FDMC method, SGBM or the CMC method.

The CMC method is a combination of the Monte-Carlo method and the COS method which can be used for computing exposures. We adapt the COS method to make it more general and thereby applicable to a wide range of possible states (S_m, v_m) , while maintaining its high accuracy. This comes at the cost of computational speed (under the Heston dynamics, particularly when the Feller condition is not satisfied). However, considering its high accuracy, this method is used as a benchmark value to analyze the accuracy and convergence of EE and PFE computed by FDMC and SGBM. By using this benchmark, it is shown that the FDMC method, SGBM and CMC method agree for multiple tests.

As a first result, it is shown that the impact of stochastic volatility on exposure profiles is most significant for $\text{PFE}_{97.5\%}$; whereas the distribution computed under

C. S. L. de Graaf et al.

the Black–Scholes dynamics suffers from the tail being chopped off at a certain maximum option value boundary, under the Heston dynamics this feature is not present.

Because any finite difference solution generates option values for an entire grid of underlying values, the FDMC method is promising. The computation time in this method is dominated by the computation of the solution on the grids at each exercise time. When these are stored, the EE computation boils down to an interpolation procedure for all paths at each exercise time. By using the COS method as a benchmark, it is shown that the error introduced by the interpolation is negligible. Compared to the CMC method, the error is within the range of 10^{-3} . A possible improvement of this method would be the implementation of a nonuniform adaptive grid, which adapts its dense region to the density of the generated paths at the specific exercise time, but this comes at a price in CPU time.

SGBM has been extended to the Heston model for computing exposures. We test the convergence of SGBM w.r.t the number of paths and the number of bundles in several ways. For the two considered tests, the computation of EE and PFE shows to be highly accurate with an error compared to CMC in the order of 10^{-3} . SGBM is an efficient Monte Carlo method when valuing exposure distributions along a time horizon.

Acknowledgments

The authors would like to thank Shashi Jain for helpful advice regarding the SGBM method and Dr. Norbert Hari for fruitful discussions about the FDMC method. Furthermore, the financial support by the Dutch Technology Foundation STW (project 12214) is gratefully acknowledged.

References

- L. Andersen (2008) Simple and efficient simulation of the Heston stochastic volatility model, *Journal of Computational Finance* **11** (3), 1–42.
- A. Antonov & D. Brecher (2012) Exposure & CVA for large portfolios of vanilla swaps: The thin-out optimization, available at SSRN 2160342.
- A. Antonov, S. Issakov & S. Mechkov (2011) Algorithmic exposure and CVA for exotic derivatives, available at SSRN 1960773, pp. 1–31.
- M. Brennan & E. Schwartz (1977) The valuation of American put options, *The Journal of Finance* **32** (2), 449–462.
- F. Fang (2010) *The COS Method: An Efficient Fourier Method for Pricing Financial Derivatives Dissertation*, Delft University of Technology, Faculty of Electrical Engineering, Mathematics and Computer Science.
- F. Fang & C. W. Oosterlee (2011) A fourier-based valuation method for Bermudan and Barrier options under Heston’s model, *SIAM Journal on Financial Mathematics* **2** (1), 439–463.
- S. Foulon & K. In’t Hout (2010) ADI finite difference schemes for option pricing, *International Journal of Numerical Analysis and Modeling* **7** (2), 303–320.
- J. Gregory (2010) *Counterparty Credit Risk*. New York: Wiley Online Library.

Efficient Computation of Exposure Profiles for CCR

- T. Haentjens & K. In't Hout (2011) ADI finite difference schemes for the Heston-Hull-White PDE, *Journal of Computational Finance* **16**(1), 83–110.
- T. Haentjens & K. J. In't Hout (2013) ADI schemes for pricing American options under the Heston model, pp. 1–26. arXiv:1309.0110.
- S. Jain & C. Oosterlee (2012) Pricing high-dimensional Bermudan options using the stochastic grid method, *International Journal of Computer Mathematics* **89**(9), 1186–1211.
- S. Jain & C. Oosterlee (2013) The stochastic grid bundling method: Efficient pricing of Bermudan options and their Greeks, pp. 1–32. Available at SSRN: <http://ssrn.com/abstract=2293942>, 1–26.
- P. Klein & J. Yang (2010) Counterparty credit risk and American options, *The Journal of Derivatives* **20**(4), 7–21.
- F. Longstaff & E. Schwartz (2001) Valuing American options by simulation: A simple least-squares approach, *Review of Financial Studies* **1** (I), 113–147.
- D. Tavella & C. Randall (2000) *Pricing Financial Instruments: The Finite Difference Method*. New York: Wiley.
- M. Vellekoop & H. Nieuwenhuis (2009) A tree-based method to price American options in the Heston model, *Journal of Computational Finance* **13**(1), 1–21.
- S. Zhu & M. Pykhtin (2007) A guide to modeling counterparty credit risk, *GARP Risk Review*, July/August, pp. 16–22.