



Research Paper

Efficient computation of exposure profiles on real-world and risk-neutral scenarios for Bermudan swaptions

Qian Feng,¹ Shashi Jain,³ Patrik Karlsson,³
Drona Kandhai^{3,4} and Cornelis W. Oosterlee^{1,2}

¹Centrum Wiskunde and Informatica (CWI), Science Park 123, 1098 XG Amsterdam, The Netherlands; emails: qian@cw.nl, c.w.oosterlee@cw.nl

²Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands

³ING Bank, Foppingadreef 7, PO Box 1800, 1000 BV Amsterdam, The Netherlands; emails: shashi.jain@ingbank.com, patrik.karlsson@ingbank.com

⁴University of Amsterdam, PO Box 94216, 1090 GE, Amsterdam, The Netherlands; email: drona.kandhai@ingbank.com

(Received June 6, 2016; accepted June 7, 2016)

ABSTRACT

This paper presents a computationally efficient technique for the computation of exposure distributions at any future time under the risk-neutral and some observed real-world probability measures; these are needed for the computation of credit valuation adjustment (CVA) and potential future exposure (PFE). In particular, we present a valuation framework for Bermudan swaptions. The essential idea is to approximate the required value function via a set of risk-neutral scenarios and use this approximated value function on the set of observed real-world scenarios. This technique significantly improves the computational efficiency by avoiding nested Monte Carlo simulation and using only basic methods such as regression. We demonstrate the benefits

of this technique by computing exposure distributions for Bermudan swaptions under the Hull–White and G2++ models.

Keywords: credit valuation adjustment (CVA); credit exposure; potential future exposure (PFE); Bermudan swaption; risk-neutral measure; real-world measure.

1 INTRODUCTION

The aim of the regulatory capital base in the Basel framework is to improve a bank's resilience against future losses due to defaults of counterparties (Basel Committee on Banking Supervision 2010). Credit exposure to counterparties occurs due to financial transactions or investments via over-the-counter (OTC) derivatives products. It is defined as the market value of the replacement costs of transactions if a counterparty defaults, assuming no recovery. Banks are required to hold regulatory capital to back exposure in the future to all their counterparties.

The Basel Committee gives specific definitions for the credit exposure metrics and adjustments regarding the future credit risk to banks/firms (Basel Committee on Banking Supervision 2005). For example, the expected exposure (EE) is the mean of the exposure distribution at any particular future date. The potential future exposure (PFE) is a high quantile (typically 97% or 99%) of the exposure distribution at any particular future date. The (unilateral) credit valuation adjustment (CVA) is the market value of the credit risk of the counterparty to the bank, which is typically calculated via an integral over time of the product of the discounted EE, the default probability and the percentage of loss given default (LGD) (Zhu and Pykhtin 2007).

EE and PFE are important indicators for the safety of a bank's portfolio to market movements. They are therefore used as metrics for capital requirements by regulators in Basel II and III (Gregory 2010). PFE is used for trading limits for portfolios with counterparties, as it may indicate at any future date the maximum amount of exposure with a predefined confidence. For example, the 99% PFE is the level of potential exposure that can be exceeded with a probability of 1%. CVA is a charge that has a direct impact on the balance sheet and the income statement of a firm, as it is an adjustment to the value of financial derivatives.

There are three basic steps in calculating future distributions of exposure (Gregory 2010):

- the generation of scenarios using the models that represent the evolution of the underlying market factors;
- the valuation of the portfolio for each scenario at each monitoring date;
- the determination of exposure values at each date for each scenario.

There is no doubt that CVA must be computed under the risk-neutral measure, as it is the market price of counterparty default risk. It is the cost of setting up a hedge portfolio to mitigate the credit risk that arises from exposure against a counterparty. In the setting of a CVA computation, scenarios are generated under the risk-neutral measure to compute “risk-neutral exposure distributions”.

In contrast, for risk analysis, it is argued that expectations (EEs) and quantiles (PFEs) of future exposure values must be obtained via scenarios that can reflect the real world in a realistic way. We know that the risk-neutral probability measure used in the pricing process does not reflect the real-world probability of future outcomes, as it has been adjusted based on the assumption that market participants are risk neutral.

The Girsanov theorem states that the risk-neutral volatility should be equal to the real-world volatility when an equivalent measure exists (Andersen and Piterbarg 2010). However, it is well known that in practice the risk-neutral market-implied volatility differs from the observed real-world volatility (Hull *et al* 2014; Stein 2013). The observed historical dynamics and the calibrated risk-neutral dynamics may exhibit a different behavior, which is a challenge for risk management, as the computational cost becomes high.

In practice, calculation of exposure values on each real-world scenario at each monitoring date needs to be performed under a risk-neutral measure. For certain products, such as Bermudan swaptions, the valuation is based on Monte Carlo simulations, which can be computationally intensive, especially since pricing then requires another nested set of Monte Carlo paths. The computational cost increases drastically due to the number of real-world scenarios, risk-neutral paths and monitoring dates.

Employing a simplification, ie, assuming that the observed real-world scenarios are close to the risk-neutral scenarios and calculation takes place under just one measure, may lead to serious problems, as there are significant differences between the resulting distributions. Stein (2014) showed that exposures computed under the risk-neutral measure depend on the choice of numéraire and can be manipulated by choosing a different numéraire. As a conclusion, it is crucial that calculations of EE and PFE are done under the real-world instead of the risk-neutral measure.

The computational problem poses a great challenge to practitioners to enhance computational efficiency. Available solutions include reduction of the number of monitoring dates and Monte Carlo paths, application of variance reduction techniques and using interpolation and enhanced computational platforms such as graphics processing units (GPUs). Even with all these efforts, calculations cost a lot of time (Stein 2014).

For Bermudan swaptions, Joshi and Kwon (2016) provided an efficient approach for approximating CVA, which relies only on an indicator of future exercise time along scenarios, the decision of which is based on the regressed functions. The expected

exposure at a monitoring date is then obtained from the corresponding deflated path-wise cashflows. However, this approximation method cannot, in a straightforward fashion, be used for PFE on the real-world scenarios. For PFE computations, Stein (2013) proposed to avoid nested Monte Carlo simulations by combining the real-world and the risk-neutral probability measures. The computed results lie between the computed PFE values under the real-world and risk-neutral probability measures.

In this paper, we will focus on accurate computation of these risk measures for a heavily traded OTC derivative, the Bermudan swaption. There are well-developed methods that can be used to compute the time-zero value of Bermudan swaptions, such as regression and simulation-based Monte Carlo methods, eg, the least squares method (LSM) (Andersen 1999; Longstaff and Schwartz 2001) or the stochastic grid bundling method (SGBM) (Jain and Oosterlee 2012, 2015; Karlsson *et al* 2014), the finite difference (FD) PDE method or the Fourier expansion-based COS method (Fang and Oosterlee 2009).

This paper presents an efficient method to significantly enhance the computational efficiency of exposure values computation without the nested simulation. The key is to approximate the value function by a linear combination of basis functions obtained by risk-neutral scenarios, and to compute the expected payoff using the approximated value function to determine the optimal early exercise strategy on the paths representing the observed real-world scenarios. Only two sets of scenarios, one under the risk-neutral and one under the observed historical dynamics, are needed to compute the exposure distributions at any future time under the two measures. We apply this numerical scheme within the context of the LSM and SGBM approaches.

The paper is organized as follows. Section 2 presents the background mathematical formulation of EE, PFE and CVA as well as the dynamic programming framework for pricing Bermudan swaptions. Section 3 explains the essential insight for computation under two measures based on the risk-neutral scenarios, and describes the algorithms for computing the exposure profiles for the SGBM and LSM. We provide reference values for exposure, based on Fourier-cosine expansions, in Section 4. Section 5 presents numerical results with the algorithms developed for the one-factor Hull–White and two-factor G2++ models.

2 CREDIT VALUATION ADJUSTMENT, EXPECTED EXPOSURE AND POTENTIAL FUTURE EXPOSURE AS RISK MEASURES

In this section, we present the general framework for computing the exposure measurements. It is important to choose suitable probability measures to compute CVA, EE and PFE. We will discuss the practical background and the choice of probability measures.

2.1 Calibration and backtesting

It is well known that there are differences between calibrated historical dynamics and the dynamics implied by market prices. The reason is that models calibrated to historical data tend to reflect future values based on historical observations, and models calibrated to market prices reflect market participants' expectations about the future. Some research on building a joint framework in the real and risk-neutral worlds is done by Hull *et al* (2014). They propose a joint measure model for the short rate, in which both historical data and market prices can be used for calibration, and the calibrated risk-neutral and real-world measures are equivalent.

The practical setting with respect to calibrating model parameters is involved, however. Backtesting of counterparty risk models is required by the Basel Committee for those banks with an internal model method approval, for which PFE is an important indicator for setting limits. Backtesting refers to comparison of the outcomes of a bank's model against realized values in the past. The bank's model must be consistent with regulatory constraints; in other words, it must be able to pass the backtesting of PFE. A bank has to strike a balance between managing its risk and meeting the expectations of the shareholders. An overconservative estimate of market factors for exposure computation would lead to high regulatory capital reservings.

In short, a model used by a bank for generating scenarios should be able to reflect the real world: it should be able to meet the requirements of backtesting limits by regulators and the return rate by investors. Based on this, Kenyon *et al* (2015) proposed a risk-appetite measure that would fit in with these requirements. When a calibrated model under this risk-appetite measure cannot pass the backtesting, the bank needs to reconsider its preferences. From backtesting, one may find a so-called PFE-limit implied volatility of a model, by which, combined with a given budget, a bank's risk preference can be computed.

Ruiz (2012) called the model that describes the evolution of the underlying market factors the risk factor evolution (RFE) model, on which the backtesting is done periodically. The related probability measure is called the RFE measure. In that work, the model used to describe the real world is introduced first, and the relevant probability measure is defined based on the model. In some sense, there are different probability measures induced by the backtesting setting that describe the outcome, assuming the underlying factors evolve according to the calibrated model.

2.2 Mathematical formulation

Consider an economy within a finite time horizon $[0, T]$. The probability space $(\Omega, \mathcal{F}, \mathbb{P})$ describes the uncertainty and information, with Ω being the sample space consisting of outcome elements ω , with \mathcal{F} being a σ -algebra on Ω , and with $\mathbb{P}: \mathcal{F} \rightarrow [0, 1]$ being the probability measure that specifies the probability of events

happening on the measure space (Ω, \mathcal{F}) . Information up to time t is included in the filtration $\{\mathcal{F}_t, t \in [0, T]\}$.

Further assume a complete market without arbitrage opportunities. There exists an equivalent martingale measure such that a price associated to any attainable claim is computed as an expectation under this probability measure with respect to the associated numéraire. We choose to use a risk-neutral probability measure, denoted by $\mathbb{Q}: \mathcal{F} \rightarrow [0, 1]$, with numéraire $B_t(w) = \exp(\int_0^t r_s(w) ds)$, where $\{r_s, s \in [0, t]\}$ is the risk-neutral short rate. The numéraire B_t represents a savings account with $B_0(w) = 1$.

Inspired by Kenyon *et al* (2015) and Ruiz (2012), we define a probability measure of observed history that can pass the backtesting. We use the notation $\mathbb{A}: \Omega' \rightarrow [0, 1]$ to present the observed historical probability measure on some measure space (Ω', \mathcal{F}') that we choose to reflect the probability of events in the real world. The probability measure $\mathbb{A}(\Omega') = 1$. The observed historical measure \mathbb{A} may not be equivalent to the chosen risk-neutral measure \mathbb{Q} . As a probability space that includes realized outcomes in the past, the observed measure space should satisfy $\Omega' \subset \Omega$ and the associated filtration $\mathcal{F}'_t \subset \mathcal{F}_t$.

Let the stochastic process $\{X_t \in \mathbb{R}^d, t \in [0, T]\}$ on (Ω, \mathcal{F}) represent all influential market factors. We further define the market factor $\{X_t\}_0^T$ on the space (Ω', \mathcal{F}') as the same mapping as the one on (Ω, \mathcal{F}) , ie, for an outcome w that may happen in both Ω and Ω' with different probability, one will have the same realized values for the market factors. Fixing an outcome $w \in \Omega' \subset \Omega$, the stochastic process is a function of time t , ie, $X_t(w): [0, T] \rightarrow \mathbb{R}^d$, which is a path of X_t .

2.3 Definition of exposure, credit valuation adjustment and potential future exposure

Let the value of a portfolio v at time t be denoted by random variable $v_t: \Omega \rightarrow \mathbb{R}$; $v_t(w)$ is the value of the portfolio at time t on a path, which is the mark-to-market value of the portfolio computed under the risk-neutral measure \mathbb{Q} .

We define exposure as the replacement costs of the portfolio, given by

$$E_t(w) = \max(0, v_t(w)), \quad (2.1)$$

where $w \in \Omega$. Once the contract expires or, in the case of early exercise options, when the contract is exercised before expiry, the exposure of the portfolio is equal to zero.

Assume the percentage of LGD to be a constant over time, and let $PS(t)$ represent the default probability up to time t , which is retrieved from credit default swap (CDS) market data under the risk-neutral probability measure. Assume the independence of

exposure and the probability of default. The CVA formula is then given by

$$\text{CVA}_0 = \text{LGD} \int_0^T \text{EE}^*(t) \, d\text{PS}(t), \quad (2.2)$$

where the notation $d\text{PS}(t)$ represents the probability that the default event occurs during the interval $[t, t + dt]$, and the discounted expected exposure EE^* is the conditional expectation of discounted exposure computed with the probability measure \mathbb{Q} , given by

$$\text{EE}^*(t) = \mathbb{E}^{\mathbb{Q}} \left[\frac{E_t}{B_t} \right] = \int_{\Omega} \frac{E_t(w)}{B_t(w)} \, d\mathbb{Q}(w), \quad (2.3)$$

where $\mathbb{E}^{\mathbb{Q}}$ is the risk-neutral expectation.

The curve $\text{PFE}(t)$ is a function of future time t until the expiry of the transactions T . Its peak value indicates the maximum potential exposure of a portfolio over the horizon $[0, T]$. We define the PFE curve at time $t \in [0, T]$ as the 99% quantile of the exposure distribution, measured by the observed probability measure \mathbb{A} , given by

$$\text{PFE}(t) = \inf\{y \mid \mathbb{A}(\{w : E_t(w) < y\}) \geq 99\%\}, \quad (2.4)$$

where $w \in \Omega'$ and $X_0(w) = x$.

The maximum PFE (MPFE) is used to measure the peak value at the PFE curve over the time horizon $[0, T]$, given by

$$\text{MPFE} = \max_{t \in [0, T]} \text{PFE}(t). \quad (2.5)$$

In a similar way, another measure of credit risk of a portfolio is the EE, which is the average exposure at any future date, denoted by $\text{EE}(t)$. The value of the EE curve at a monitoring date t under the observed measure \mathbb{A} is given by

$$\text{EE}(t) = \mathbb{E}^{\mathbb{A}}[E_t] = \int_{\Omega'} E_t(w) \, d\mathbb{A}(w), \quad (2.6)$$

where $w \in \Omega'$ and $X_0(w) = x$. The real-world expected positive exposure (EPE) over a time period $[0, T]$ is given by

$$\text{EPE}(0, T) = \frac{1}{T} \int_0^T \text{EE}(t) \, ds. \quad (2.7)$$

In particular, we are interested in Bermudan swaptions, the pricing dynamics of which are presented in the following section.

2.4 Pricing of Bermudan swaptions

A Bermudan swaption is an option where the owner has the right to enter into an underlying swap either on the swaption's expiry or at a number of other predefined exercise dates before the expiry date. As soon as the swaption is exercised, the underlying swap starts. We assume here that the expiry date of the swap is predefined, so the duration of the swap is calculated from the swaption exercise date until a fixed end date. The underlying dynamics for the short rate governing the Bermudan swaption are either the one-factor Hull–White model or the two-factor G2++ model. Details of these well-known governing dynamics, either under the risk-neutral or the observed real-world dynamics, are presented in Appendixes 1 and 2 (available online).

We assume that the exercise dates coincide with the payment dates of the underlying swaps. Then, we consider an increasing maturity structure, $0 < T_1 < \dots < T_N < T_{N+1}$, with T_{N+1} the fixed end date of the underlying swap and T_1, T_N the first and last opportunities to enter, respectively. We define $T_0 = 0$. We assume that when an investor enters a swap at time T_n , $n = 1, 2, \dots, N$, the payments of the underlying swap will occur at $T_{n+1}, T_{n+2}, \dots, T_{N+1}$, with time fraction $\tau_n = T_{n+1} - T_n$. We let N_0 represent the notional amount and K be the fixed strike. We use indicator $\delta = 1$ for a payer Bermudan swaption and $\delta = -1$ for a receiver Bermudan swaption.

The payoff for entering the underlying swap at time T_n associated with payment times $\mathcal{T}_n = \{T_{n+1}, \dots, T_{N+1}\}$, conditional on $X_{T_n} = x$, is given by (Brigo and Mercurio 2007)


$$U_n(x) = N_0 \left(\sum_{k=n}^N P(T_n, T_{k+1}, x) \tau_k \right) \max(\delta(S(T_n, \mathcal{T}_n, x) - K), 0), \quad (2.8)$$

where the forward swap rate $S(t, \mathcal{T}_n, x)$ at time $t \leq T_n$ associated with time T_n, \dots, T_{N+1} is defined by

$$S(T_n, \mathcal{T}_n, x) = \frac{1 - P(T_n, T_{N+1}, x)}{\sum_{k=n}^N P(T_n, T_{k+1}, x) \tau_k}, \quad (2.9)$$

and $P(T_n, T_k, x)$ is the price of a zero-coupon bond (ZCB), conditional on $X_{T_n} = x$, associated with times T_n and T_k . The analytic formula of the ZCB is related to the risk-neutral model for the underlying variable (see, for example, Appendixes 1 and 2, available online).

We refer to a function U_n , a bounded Borel function, as the exercise function, which represents the value of the future payments on any given scenario, when the option will be exercised at time T_n . For completeness, we define $U_0 \equiv 0$. We choose for the stochastic process $\{X_t, t \in [0, T_N]\}$ an Ito diffusion. In that case, $U_n(X_{T_n})$

FIGURE 1 Time lines.

is a continuous variable, as X_{T_n} is a continuous random variable. The value of not exercising the option at $t \in [0, T_N)$ is the value of continuing the option at time t .

Let time $t \in [T_n, T_{n+1})$, where the exercise opportunities are restricted to dates $\{T_{n+1}, \dots, T_N\}$. The value of the Bermudan claim is the risk-neutral expectation of the (discounted) future payoff when exercising optimally (Øksendal 2003). With the strong Markov property of the Ito diffusions (Øksendal 2003), the value of this Bermudan claim at time t , conditional on $X_t = x$, is the value that is obtained by maximizing the following object function (Glasserman 2003):

$$C(t, x) = \max_{I \in \{n+1, \dots, N\}} B_t \mathbb{E}^{\mathbb{Q}} \left[\frac{U_I(X_{T_I})}{B_{T_I}} \mid X_t = x \right], \quad (2.10)$$

where $n = 0, \dots, N - 1$. We refer to the value function $C(t, \cdot)$ as the continuation function at time t .

We wish to determine the exposure at a set of discrete monitoring dates, $\{0 = t_0 < t_1 < \dots < t_M = T_N\}$, with time step $\Delta t_k = t_{k+1} - t_k$, $k = 0, \dots, M - 1$. These monitoring dates include the exercise dates $\{T_1, T_2, \dots, T_N\}$, and t_M is equal to T_N . There are some dates between each two exercise dates, as we are also interested in the exposure at those intermediate dates.

Figure 1 presents the time lines of the exercise dates of a Bermudan swaption and the monitoring dates used for exposure computation as an example. This Bermudan swaption can be exercised seven times between year four and year ten, ie, year four is the first exercise date and year ten is the expiry (the last exercise date). The exposure monitoring dates are each one-fifth of a year from time zero until year ten. The monitoring date $t_{20} = 4$ coincides with the first exercise date, and the monitoring date $t_{50} = 10$ is equal to the last exercise opportunity.

We compute the exposure of a Bermudan claim at monitoring dates $\{t_m\}_{m=0}^M$. Value function V then satisfies (Glasserman 2003)

$$V(t_m, x) = \begin{cases} U_N(x), & t_M = T_N, \\ \max(C(t_m, x), U_n(x)), & t_m = T_n, n < N, \\ C(t_m, x), & T_n < t_m < T_{n+1}, n < N, \end{cases} \quad (2.11)$$

where the continuation function C is computed as the conditional expectation of the future option value, given by

$$C(t_m, x) = B_{t_m} \mathbb{E}^{\mathbb{Q}} \left[\frac{V(t_{m+1}, X_{t_{m+1}})}{B_{t_{m+1}}} \mid X_{t_m} = x \right], \quad (2.12)$$

which can be proven to be equivalent to (2.10) by induction.

The optimal exercise strategy is now as follows. At state $X_{T_n} = x$, exercise takes place when $U_n(x) > C(T_n, x)$, and the option is kept at all non-exercise monitoring dates t_m . The value function V and continuation function C are defined over the time period $[0, T_N]$ and space $\mathcal{D} \in \mathbb{R}^d$.

The pricing dynamics in (2.11) are most conveniently handled by means of a backward recursive iteration. From known value U_N at time $t_M = T_N$, we compute $V(t_{M-1}, \cdot)$, and subsequently function $V(t_{M-2}, \cdot)$, and so on, until time zero. The essential problem, hence, becomes to determine the value function V and continuation function C at all monitoring dates $\{t_m\}_{m=1}^M$.

REMARK 2.1 Given a fixed path $w' \in \Omega'$ or $w \in \Omega$, we compute the option values for the scenario as $V(t_m, X_{t_m}(w))$ at any monitoring date t_m by (2.11). Once the option for scenario w is exercised at a specific date, the option terminates, and the exposure values regarding this option along the scenario from the exercise date to T become zero.

When a sufficient number of scenarios for the risk-neutral model are generated, the option value can be determined at all monitoring dates for any scenario, and we obtain a matrix of exposure values called the exposure profile.

The exposure profile, computed from observed real-world scenarios that are calibrated based on historical data, is an empirical real-world exposure density from which we can estimate real-world EEs and PFEs at each monitoring date. However, with risk-neutral short rate processes, the exposure profiles on risk-neutral scenarios are needed to compute the discounted EE.

We see that the key to computing exposure profiles on generated scenarios is to know the value function V and the continuation function C at all monitoring dates $\{t_m\}_{m=1}^M$.

Nested Monte Carlo simulation is often used when a simulation-based algorithm is employed for the valuation; this is expensive, as simulations of risk-neutral paths are

needed for each (real-world) scenario at each monitoring date. Suppose that accurate pricing requires K_I risk-neutral paths at M monitoring dates; then, the computational time would be $O(M^2 K_I K_a)$ for K_a real-world scenarios for computing EE and PFE profiles.

3 REGRESSION-BASED MONTE CARLO ALGORITHMS

The computation of the conditional risk-neutral expectation is the most expensive part in the algorithm for dynamics (2.11). We propose algorithms that can approximate the continuation function in (2.12) by basic functions (for example, polynomial functions), based on the risk-neutral scenarios. Using these functions, we can perform simulations with risk-neutral expectations on the real-world scenarios without nested simulations. To compute CVA and PFE, we only need one set of K_q risk-neutral scenarios and one set of K_a real-world scenarios.

The proposed algorithms are based on the approximation of the continuation function within the SGBM and LSM simulation techniques. In this section, details of the algorithms are presented as well as the differences between the LSM and SGBM.

3.1 Stochastic grid bundling method

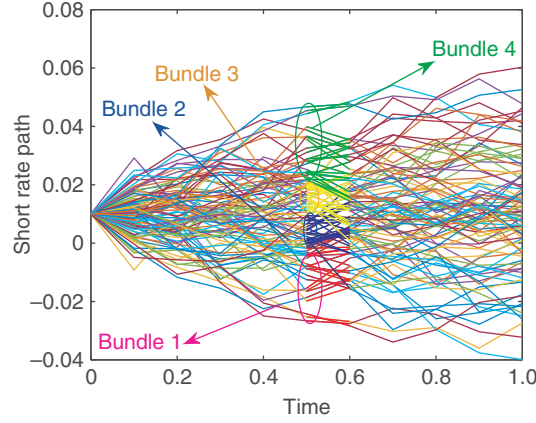
The SGBM approach, based on regression, bundling and simulation, was developed by Jain and Oosterlee (2015) for pricing Bermudan options. The SGBM can be very naturally generalized toward the efficient computation of exposure profiles because of its high accuracy in approximating expected payoffs on each Monte Carlo path. The SGBM has been used to compute risk-neutral exposure profiles (for computing CVA) of Bermudan-style claims in Karlsson *et al* (2014) and Feng and Oosterlee (2014).

Pricing in the context of the SGBM approach is based on risk-neutral scenarios. Computation of discounted expected option values is performed locally in so-called bundles by means of local regression. We will store the bundle-wise approximated continuation functions and use them to compute exposure profiles for the observed real-world scenarios for a Bermudan swaption.

3.1.1 Risk-neutral scenarios

Let $\{X_{1,h}^q, \dots, X_{M,h}^q\}_{h=1}^{K_q}$ be K_q scenarios, where the underlying factor evolves with the risk-neutral model. Pricing is done by a backward-in-time iteration, as in (2.11), from time t_M to time $t_0 = 0$.

To initialize the computation, the option value at expiry $t_M = T_N$ is computed as the immediate payoff U_N , ie, the option value realized on the h th scenario at time

FIGURE 2 Bundles and the disjoint sub-domains at time 0.5.

t_M , $\hat{v}_{M,h}^q = U_N(X_{M,h}^q)$. As the option either expires or is exercised at time t_M , the exposure equals zero for all paths at time t_M , $\{\hat{E}_{M,h}^q = 0\}_{h=1}^{K_q}$.

At monitoring dates t_m , $m = M - 1, \dots, 1$, J partitions $\{\mathcal{B}_{m,j}\}_{j=1}^J$, which are called bundles, are defined; these consist of Monte Carlo path values at t_m and have very similar realized values based on the cross-sectional risk-neutral samples $\{X_{m,h}^q\}_{h=1}^{K_q}$. The realized values of the risk-neutral paths form a bounded domain, and these bundles divide the domain into disjoint sub-domains $\{\mathcal{D}_{m,j}\}_{j=1}^J$. For a one-dimensional variable, these disjoint sub-domains connected to bundles can be presented by

$$\mathcal{D}_{m,j} = \left(\max_{h \in \mathcal{B}_{m,j-1}} (X_{m,h}^q), \max_{h \in \mathcal{B}_{m,j}} (X_{m,h}^q) \right], \quad (3.1)$$

where $j = 2, 3, \dots, J - 1$. In particular, we define the first sub-domain at the left-side boundary as

$$\mathcal{D}_{m,1} = \left(\min_{h \in \mathcal{B}_{m,1}} (X_{m,h}^q), \max_{h \in \mathcal{B}_{m,1}} (X_{m,h}^q) \right]. \quad (3.2)$$

Figure 2 shows the Monte Carlo paths in four bundles and the associated disjoint sub-domains at a monitoring date. The bundles are based on the values of the realized short rate at time 0.5.

The continuation function $C(t_m, \cdot)$ is approximated “in a bundle-wise fashion” for each domain by approximating the value function at time t_{m+1} for the paths in a bundle.

For $j = 1, \dots, J$, on the Monte Carlo paths in bundle $\mathcal{B}_{m,j}$, the value function $V(t_{m+1}, \cdot)$ is approximated by a linear combination of basis functions $\{\phi_k\}_{k=1}^B$, ie,

$$V(t_{m+1}, y) \approx \sum_{k=1}^B \beta_k(m, j) \phi_k(y), \quad (3.3)$$

where the coefficients $\beta_k(m, j)$ of the k th basis function minimize the sum of squared residuals over the paths in bundle $\mathcal{B}_{m,j}$, ie,

$$\sum_{h \in \mathcal{B}_{m,j}} \left(\hat{v}_{m+1,h}^q - \sum_{k=1}^B \beta_k(m, j) \phi_k(X_{m+1,h}^q) \right)^2, \quad (3.4)$$

with $\{\hat{v}_{m+1,h}^q\}_{h=1}^{K_q}$ the option values at time t_{m+1} on the cross-sectional sample $\{X_{m+1,h}^q\}_{h=1}^{K_q}$.

Using the approximated value function in (3.3) instead of the “true value” in (2.12), the continuation function on $\mathcal{D}_{m,j}$ can be approximated by

$$C(t_m, x) \approx \sum_{k=1}^B \beta_k(m, j) \psi_k(x, t_m, t_{m+1}), \quad (3.5)$$

where $x \in \mathcal{D}_{m,j}$ and function ψ_k is the conditional risk-neutral discounted expectation of basis function ϕ_k , defined by

$$\psi_k(x, t_m, t_{m+1}) := B_{t_m} \mathbb{E}^{\mathbb{Q}} \left[\frac{\phi_k(X_{t_{m+1}})}{B_{t_{m+1}}} \mid X_{t_m} = x \right]. \quad (3.6)$$

The formulas for $\{\psi_k\}_{k=1}^B$ can be obtained easily, and often analytically, when polynomial terms are chosen as the basis functions (see Section 3.1.4).

The expected values on the paths of the bundle $\mathcal{B}_{m,j}$ can then be approximated by

$$\hat{c}_{m,h}^q \approx \sum_{k=1}^B \beta_k(m, j) \psi_k(X_{m,h}^q, t_m, t_{m+1}), \quad (3.7)$$

where $h \in \mathcal{B}_{m,j}$.

After computation of the continuation values for all paths $\{\hat{c}_{m,h}^q\}_{h=1}^{K_q}$ at time t_m , we determine the option value at time t_m by

$$\hat{v}_{m,h}^q = \begin{cases} \max(U_n(X_{m,h}^q), \hat{c}_{m,h}^q), & t_m = T_n, \\ \hat{c}_{m,h}^q, & t_m \in (T_n, T_{n+1}), \end{cases} \quad (3.8)$$

where U_n is the exercise function.

The exposure value on the h th path from time t_m to expiry t_M is updated by the following scheme.

- (1) When exercised at exercise time $t_m = T_n$, a value of zero is assigned to the exposures along the path from time t_m to expiry, ie,

$$\hat{E}_{k,h}^q = 0, \quad k = m, \dots, M.$$

- (2) When the option is “alive” at an exercise date, or when t_m is a monitoring date between two exercise dates, the exposure at the path is equal to the approximated continuation value, $\hat{E}_{m,h}^q = \hat{c}_{m,h}^q$, and the exposure values at later times remain unchanged.

The algorithm proceeds by moving one time step backward to t_{m-1} , where the paths are again divided into new bundles, based on the realized values $\{X_{m-1,h}^q\}_{h=1}^{K_q}$, and the continuation function is approximated in a bundle-wise fashion. Option values are evaluated, and the exposure profile is updated. The algorithm proceeds, recursively, back to $t_0 = 0$. At time t_0 , we do not need bundles, and regression takes place for all paths to get the coefficients $\{\beta_k(0)\}_{k=1}^B$, ie, the option value at time zero is approximated by

$$\hat{v}_0^q \approx \sum_{k=1}^B \beta_k(0) \psi_k(x_0, t_0, t_1). \quad (3.9)$$

During the backward recursive iteration, information about the boundaries of the disjoint sub-domains, $\mathcal{D}_{m,j}$, is stored, along with the associated coefficients $\{\beta_k(m, j)\}_{k=1}^B$ for each index, $j = 1, \dots, J$, at each monitoring date, t_m , $m = 0, \dots, M-1$. Based on this information, we can retrieve the piecewise approximated continuation function for each time t_m .

With the risk-neutral exposure profiles, $\{\hat{E}_{1,h}^q, \dots, \hat{E}_{M,h}^q\}_{h=1}^{K_q}$, the discounted EE of a Bermudan swaption can be approximated by

$$\text{EE}^*(t_m) \approx \frac{1}{K_q} \sum_{h=1}^{K_q} \exp\left(-\sum_{k=0}^{m+1} \frac{1}{2}(\hat{r}_{k,h}^q + \hat{r}_{k+1,h}^q)\Delta t_k\right) \hat{E}_{m,h}^q, \quad (3.10)$$

where $\{\hat{r}_{1,j}^q, \dots, \hat{r}_{M,j}^q\}_{j=1}^{K_q}$ represents simulated risk-neutral short rate values.

3.1.2 Real-world scenarios

During the computations on the risk-neutral scenarios, we have stored the bundle-wise coefficients $\{\beta_k(m, j)\}_{k=1}^B$ and the associated sub-domains $\{\mathcal{D}_{m,j}\}_{j=1}^J$, by which we can perform valuation and exposure computation for any scenario without nested simulation.

We present the steps to compute exposure profiles on a set of K_a observed real-world scenarios $\{X_{1,h}^a, \dots, X_{M,h}^a\}$. These profiles are also determined by a backward iteration from time t_M until time t_0 .

At expiry date t_M , the exposure equals zero,

$$\{\hat{E}_{M,h}^a = 0\}_{h=1}^{K_a}.$$

At monitoring dates $t_m < t_M$, for each index $j = 1, \dots, J$, we determine those paths for which $X_{m,h}^a \in \mathcal{D}_{m,h}$; we compute the continuation values for these paths by

$$\hat{c}_{m,h}^a \approx \sum_{k=1}^B \beta_k(m, j) \psi_k(X_{m,h}^a, t_m, t_{m+1}), \quad (3.11)$$

where $X_{m,h}^a \in \mathcal{D}_{m,j}$.

Based on these continuation values, we update the exposure profile on this set of real-world scenarios.

At an exercise time $t_m = T_n$, we compare the approximated continuation value $\hat{c}_{m,h}^a$ with the immediate exercise values $U_n(X_{m,h}^a)$ for each path; when the immediate exercise value is largest, the option is exercised at this path at time t_m , and exposure values at this path from time t_m to expiry are set to zero, ie, $\hat{E}_{k,h}^a = 0$, $k = m, \dots, M$.

Otherwise, $\hat{E}_{m,h}^a = \hat{c}_{m,h}^a$ and the later exposure values remain unchanged.

When t_m is an intermediate monitoring date, the exposure values are equal to the continuation values in (3.11).

Note that the time-zero option value is the same for the risk-neutral and real-world scenarios, ie, $\hat{v}_0^q = \hat{v}_0^a$. Values of the observed real-world PFE and EE curves at monitoring dates t_m can be approximated by

$$\begin{aligned} \text{PFE}(t_m) &= \text{quantile}(\hat{E}_{m,h}^a, 99\%), \\ \text{EE}(t_m) &= \frac{1}{K_a} \sum_{h=1}^{K_a} \hat{E}_{m,h}^a. \end{aligned} \quad (3.12)$$

3.1.3 Stochastic grid bundling method bundling technique

An essential technique within the SGBM is the bundling of asset path values at each monitoring date, based on the cross-sectional risk-neutral samples. Numerical experiments have shown that the algorithm converges with respect to the number of bundles (Feng and Oosterlee 2014; Jain and Oosterlee 2015).

Various bundling techniques have been presented in the literature, such as the recursive-bifurcation method, k -means clustering (Jain and Oosterlee 2015) and the equal-number bundling method (Feng and Oosterlee 2014). Here, we use the

equal-number bundling technique. In this method, at each time step t_m , we rank the paths by their realized values, $\{X_{m,h}^q\}_{h=1}^{K_q}$, and place the paths with indexes between $(j-1)K_q/J + 1$ and jK_q/J into the j th bundle, $\mathcal{B}_{m,j}$, $j = 1, \dots, J-1$. The remaining paths are placed in the J th bundle, $\mathcal{B}_{m,J}$. Asset paths do not overlap among bundles at time t_m , and each path is placed in a bundle.

The advantage of the equal-number bundling technique is that the number of paths within each bundle is proportional to the total number of asset paths. An appropriate number of paths in each bundle is important for accuracy during the local regression. As mentioned, the bundling technique is also used to determine the disjoint sub-domains on which the value function is approximated in a piece-wise fashion.

For high-dimensional problems, one can either use the equal-number bundling technique along each dimension, as employed in Feng and Oosterlee (2014), or one can project the high-dimensional vector onto a one-dimensional vector and then apply the equal-number bundling technique (see Jain and Oosterlee 2015; Leitao and Oosterlee 2015).

3.1.4 Formulas for the discounted moments in the stochastic grid bundling method

When we choose monomials as the basis functions within the bundles in the SGBM, the conditional expectation of the discounted basis functions is equal to the discounted moments. There is a direct link between the discounted moments and the discounted characteristic function (dChF), which we can also use to derive analytic formulas for the discounted moments.

As a one-dimensional example, in which the underlying variable represents the short rate, ie, $X_t = r_t$, let the basis functions be $\phi_k(r_t) = (r_t)^{k-1}$, $k = 1, \dots, B$. The discounted moments ψ_k , conditional on $r_{t_m} = x$ over the period (t_m, t_{m+1}) , are given by

$$\psi_k(x, t_m, t_{m+1}) := \mathbb{E}^{\mathbb{Q}} \left[\exp \left(- \int_{t_m}^{t_{m+1}} r_s ds \right) (r_{t_{m+1}})^{k-1} \mid r_{t_m} = x \right], \quad (3.13)$$

and the associated dChF is given by

$$\Phi(u; x, t_m, t_{m+1}) := \mathbb{E}^{\mathbb{Q}} \left[\exp \left(- \int_{t_m}^{t_{m+1}} r_s ds + iur_{t_{m+1}} \right) \mid r_{t_m} = x \right]. \quad (3.14)$$

When an explicit formula for the dChF is available, ψ_k can be derived by

$$\psi_k(x, t_m, t_{m+1}) = \frac{1}{(i)^{k-1}} \frac{\partial^{k-1} \Phi}{\partial u^{k-1}}(u; x, t_m, t_{m+1}) \Big|_{u=0}. \quad (3.15)$$

Using the relation in (3.15), we find analytic formulas for the discounted moments when the dChF is known. The dChFs of the Hull–White and G2++ models are presented in Appendixes 1 and 2, respectively (available online).

3.2 Least squares method

The LSM is also a regression-based Monte Carlo method that is very popular among practitioners. The objective of the LSM algorithm is to find for each path the optimal stopping policy at each exercise time T_n ; the option value is computed as the average value of the generated discounted cashflows. The optimal early exercise policy for the in-the-money paths is determined by comparing the immediate exercise value and the approximated continuation value, which is approximated by a linear combination of (global) basis functions $\{\phi_k\}_{k=1}^B$.

One can always combine the (expensive) nested Monte Carlo simulation with the LSM for the computation of EE and PFE on observed real-world scenarios. We will adapt the original LSM algorithm to obtain a more efficient method for computing risk-neutral and real-world exposures. The technique is similar to that described for the SGBM: valuation on the risk-neutral scenarios, approximation of the continuation function and computation of risk-neutral and real-world exposure quantities.

The involved part in the LSM is that discounted cashflows, realized on a path, are not representative of the “true” continuation values. In the LSM algorithm, the approximated continuation values are only used to determine the exercise policy; therefore, one cannot use them to determine the maximum of the immediate exercise value and discounted cashflows in order to approximate the option value (Feng and Oosterlee 2014), as is done in the SGBM.

The challenge is to approximate exposure values by means of the realized discounted cashflows over all paths.

Joshi and Kwon (2016) present a way of employing realized discounted cashflows and the sign of the regressed values for an efficient computation of CVA on risk-neutral scenarios. However, since the average of discounted cashflows is not the value of a contract under the observed real-world measure, it cannot be used to compute real-world EE or PFE quantities.

Here, we propose two LSM-based algorithms for the approximation of continuation values with realized cashflows. They can be seen as alternative algorithms to the SGBM for the computation of exposure values when we do not have expressions for the discounted moments (or when the LSM is the method of choice for many other tasks). We will test the accuracy of the algorithms compared with the SGBM and reference values generated by the COS method in Section 5.

3.2.1 Risk-neutral scenarios

First of all, we briefly explain the original LSM algorithm with the risk-neutral scenarios. At the final exercise date, $t_M = T_N$, the option holder can either exercise an option or not, and the generated cashflows are given by $q_{M,h} = U_N(X_{M,h}^q)$, $h = 1, \dots, K^q$.

At monitoring dates $t_m \in (T_{n-1}, T_n)$, at which the option cannot be exercised, the realized discounted cashflows are updated by

$$q_{m,h} = q_{m+1,h} D_{m,h}, \quad (3.16)$$

with the discount factor $D_{m,h} = \exp(-\frac{1}{2}(\hat{r}_{m,h}^q + \hat{r}_{m+1,h}^q)\Delta t_m)$.

At an exercise date $t_m = T_n$, prior to the last exercise opportunity, the exercise decision is based on the comparison of the immediate payoff by exercising and the continuation value when holding the option on the in-the-money paths; the continuation values at those in-the-money paths are approximated by projecting the (discounted) cashflows of these paths onto some global basis functions $\{\phi_1, \dots, \phi_B\}$.

The option is exercised at an in-the-money path, where the payoff is larger than the continuation value. After determining the exercise strategy at each path, the discounted cashflows read

$$q_{m,h} = \begin{cases} U_n(X_{m,h}^q), & \text{exercised,} \\ q_{m+1,h} D_{m,h}, & \text{to be continued.} \end{cases} \quad (3.17)$$

Again, computation of the discounted cashflows at any monitoring date takes place recursively, backward in time. At time $t_0 = 0$, the option value is approximated by

$$\hat{v}_{0,h}^q \approx \frac{1}{K^q} \sum_{h=1}^{K^q} q_{0,h}.$$

During the backward recursion, the discounted cashflows realized on all paths at each monitoring date t_m are computed.

For the computation of the real-world EE and PFE quantities, valuation needs to be done on the whole domain of realized asset values, as we need the continuation values at each monitoring date for all paths. We therefore propose to use the realized discounted cashflows determined by (3.17) or (3.16) on the risk-neutral scenarios.

One possible algorithm in the LSM context involves employing two disjoint sub-domains, similar to in the SGBM. At each monitoring date $t_m \in (T_{n-1}, T_n]$, Monte Carlo paths are divided into two bundles based on the realized values of the underlying variable, so the approximation can take place in two disjoint sub-domains, given by

$$\mathcal{U}_{n,1} = \{x \mid U_n(x) \leq 0\}, \quad \mathcal{U}_{n,2} = \{x \mid U_n(x) > 0\}. \quad (3.18)$$

The continuation function is approximated on these two sub-domains as

$$C(t_m, x) \approx \sum_{k=1}^B \zeta_k(t_m, j) \phi_k(x), \quad (3.19)$$

where $x \in \mathcal{U}_{n,j}$; the coefficients $\zeta_k(t_m, j)$ are obtained by minimizing the sum of squared residuals over the two bundles, respectively, which is given by

$$\sum_{X_{m,h}^q \in \mathcal{U}_{n,j}} \left(q_{m+1,h} D_{m,h} - \sum_{k=1}^B \zeta_k(t_m, j) \phi_k(X_{m,h}^q) \right)^2. \quad (3.20)$$

We refer to this technique as the LSM-bundle technique.

The other possible algorithm is to perform the regression over all Monte Carlo paths and compute the approximated continuation function on each path. The regression is as in (3.20), using basis functions and discounted cashflows but for all paths. We call this the LSM-all algorithm. Note that the exercise decision is still based on the in-the-money paths with approximated payoff, using (3.19) at exercise dates $T_n < T_N$, $n = 1, \dots, N - 1$.

We compute the risk-neutral exposure profiles with the approximated value functions in (3.19) by means of the same backward recursion procedure in Section 3.1.1.

3.2.2 Real-world scenarios

The LSM-bundle algorithm can be used for computing exposure on the observed real-world scenarios directly. It is based on the same backward iteration in Section 3.1.2; however, the continuation values are computed by the function in (3.19), ie,

$$\hat{c}_{m,h}^a \approx \sum_{k=1}^B \zeta_k(t_m, j) \phi_k(X_{m,h}^a). \quad (3.21)$$

At an early exercise date $t_m = T_n < T_N$, the early exercise policy is determined for in-the-money paths by comparing $\hat{c}_{m,h}^a$ with the immediate exercise value $U_n(X_{m,h}^a)$. Exposure values along the path from time t_m to expiry are set to zero if the option at a path is exercised.

By the LSM-all algorithm, we use the continuation function approximated in (3.19) for determining the optimal early exercise time on each real-world path; the regressed function is based on all paths to compute exposure values. We will compare the LSM-bundle and LSM-all algorithms in Section 5.

3.3 Differences between the stochastic grid bundling method and least squares method algorithms

The SGBM differs from the LSM with respect to the bundling and the local regression based on the discounted moments. By these components, the SGBM approximates the continuation function in a more accurate way than the LSM, but at a (small) additional computational cost. Here, we give some insights into these differences.

The use of SGBM bundles may improve the local approximation on the disjoint sub-domains, and we can reduce the number of basis functions.

Another important feature of the SGBM is that option values are obtained from regression in order to obtain the coefficients for the continuation function.

Loosely speaking, the continuation function is approximated locally on the bounded sub-domains $\{\mathcal{D}_{m,j}\}_{j=1}^J$ by projection on the functions $\{\psi_k\}_{k=1}^B$.

Compared with the SGBM, the LSM is based on the discounted cashflows for regression to approximate the expected payoff; however, discounted cashflows do not represent the realized expected payoff on all Monte Carlo paths. In the LSM, the expected payoff is only used to determine the optimal early exercise time and not the option value. One cannot compute the option value by using the maximum of the expected payoff and the exercise value, as it will lead to an upward bias for the time-zero option value (Longstaff and Schwartz 2001).

The SGBM does not suffer from this, and the maximum of the exercise value and the regressed continuation values gives us the direct estimator. We recommend also computing the path estimator for convergence of the SGBM algorithm. Based on a new set of scenarios with the obtained coefficients to determine the optimal exercise policy on each path, we then take the average of the discounted cashflows as the time-zero option value. Upon convergence, the direct and path estimators should be very close (Jain and Oosterlee 2015).

The LSM approach is a very efficient and adaptive algorithm for computing option values at time zero. The LSM-based algorithms for computing exposure can be regarded as alternative ways of computing the future exposure distributions based on simulation. We will analyze the accuracy of all variants in Section 5.

4 THE COS METHOD

In this section, we explain the computation of the continuation function of Bermudan swaptions under the one-factor Hull–White model by the COS method. The COS method is an efficient and accurate method based on Fourier-cosine expansions. It can be used to determine reference values for the exposure. For Lévy processes and early exercise options, the computational speed of the COS method can be enhanced by incorporating the fast Fourier transform (FFT) into the computations. We cannot employ the FFT, because the resulting matrixes with the Hull–White model do not have the special form needed (Toeplitz and Hankel matrixes (see Fang and Oosterlee 2009)) to employ the FFT. For the G2++ model, the two-dimensional COS method developed in Ruijter and Oosterlee (2012) may be used for pricing Bermudan swaptions, but this is not pursued here.

When the short rate is a stochastic process, the discount factor is a random variable, which should be under the expectation operator when computing the continuation

values. In order to compute the discounted expectation of the future option values, we will work with the discounted density function. Let $p(y, z; t, T, x)$ be the joint density function of the underlying variables $X_T = y$ and $z = -\log(B_t/B_T)$, conditional on $X_t = x$. The discounted density function is defined as the marginal probability function p_X of X_T , derived by integrating the joint density p over $z \in \mathbb{R}$:

$$p_X(y; t, T, x) := \int_{\mathbb{R}} e^{-z} p(y, z; t, T, x) dz. \quad (4.1)$$

The dChF is the Fourier transform of the discounted density function, ie,

$$\begin{aligned} \Phi(u; t, T, x) &= \mathbb{E} \left[\frac{B_t}{B_T} \exp(iuX_T) \mid X_t = x \right] \\ &= \int_{\mathbb{R}^n} e^{iu^y} \int_{\mathbb{R}} e^{-z} f(y, z; t, T, x) dz dy \\ &= \int_{\mathbb{R}^n} e^{iu^y} p_X(y; t, T, x) dy, \end{aligned} \quad (4.2)$$

where $u \in \mathbb{R}^d$ and $X_t \in \mathbb{R}^d$.

In the one-dimensional setting, $X_t = X_t$, the discounted density function p_X can be approximated by Fourier-cosine expansions (Fang and Oosterlee 2009). On an integration range $[a, b]$, we define $\{u_k\}_{k=0}^{Q-1}$ by

$$u_k = \frac{k\pi}{b-a}, \quad k = 0, \dots, Q-1,$$

where Q represents the number of cosine terms used in the Fourier-cosine expansion of the discounted density, which is given by

$$p_X(y; t, T, x) \approx \frac{2}{b-a} \sum_{k=0}^{Q-1} \mathcal{P}_k(x, t, T) \cos(u_k(y-a)). \quad (4.3)$$

The symbol \sum' in (4.3) implies that the first term of the summation is multiplied by $\frac{1}{2}$ and the Fourier coefficients \mathcal{P}_k are given by

$$\mathcal{P}_k(x, t, T) := \text{Re}\{\Phi(u_k; t, T, x) \exp(-iau_k)\}. \quad (4.4)$$

The integration range should be chosen such that the integral of the discounted density function over the region $[a, b]$ resembles very well the value of a ZCB between time t and T , given $X_t = x$. The way of constructing the range is presented in Appendix 1 (available online).

With (2.12), the continuation function, conditional on $X_{t_m} = x$, at any monitoring date $t_m \in [0, T_N)$ can be computed as an integral over $[a, b]$:

$$\begin{aligned} C(t_m, x) &\approx \int_a^b V(t_{m+1}, y) p_X(y; t_m, t_{m+1}, x) dy \\ &\approx \frac{2}{b-a} \sum_{k=0}^{Q-1} \mathcal{P}_k(x, t_m, t_{m+1}) \mathcal{V}_k(t_{m+1}), \end{aligned} \quad (4.5)$$

where the coefficients $\mathcal{V}_k(t_{m+1})$ are defined by

$$\mathcal{V}_k(t_{m+1}) := \int_a^b V(t_{m+1}, y) \cos(u_k(y-a)) dy. \quad (4.6)$$

The coefficients $\{\mathcal{V}_k\}_{k=0}^{Q-1}$ can be computed at monitoring dates $\{t_m\}_{m=1}^M$ by the backward recursion, as in (2.11). Analytic formulas in the case of the Hull–White model for the coefficients $\{\mathcal{V}_k\}_{k=0}^{Q-1}$ can be computed by backward recursion.

At the expiry date, $t_M = T_N$, the option value equals the payoff of the underlying swap, ie, $V(t_M, \cdot) = U_N(\cdot)$. We are only interested in the in-the-money region regarding the function U_n , for which we need to solve $U_N(x^*(T_N)) = 0$.

Function U_N is positive on the range $(a, x^*(T_N))$ for a receiver Bermudan swaption, and on the range $(x^*(T_N), b)$ for a payer Bermudan swaption. We compute the integral on the range in which $U_N > 0$ for the coefficients $\mathcal{V}_k(t_M)$. The formulas for the integral are given by

$$\begin{aligned} \mathcal{V}_k(t_M) &= \int_a^b U_N(y) \cos(u_k(y-a)) dy \\ &= \begin{cases} \mathcal{G}_k(a, x^*(T_N), T_N) & \text{for a receiver swaption,} \\ \mathcal{G}_k(x^*(T_N), b, T_N) & \text{for a payer swaption.} \end{cases} \end{aligned} \quad (4.7)$$

The coefficients \mathcal{G}_k at time T_n over $[x_1, x_2]$ are computed by

$$\begin{aligned} \mathcal{G}_k(x_1, x_2, T_n) &= N_0 \int_{x_1}^{x_2} \cos(u_k(y-a)) U_N(y) dy \\ &= N_0 \delta(\mathcal{A}_k^1(x_1, x_2) - \mathcal{A}_k^2(x_1, x_2, T_n)), \end{aligned} \quad (4.8)$$

where the coefficients are given by

$$\begin{aligned} &\mathcal{A}_k^1(x_1, x_2) \\ &= \begin{cases} x_2 - x_1, & k = 0, \\ \frac{1}{u_k} [\sin(u_k(x_2 - a)) - \sin(u_k(x_1 - a))], & k \neq 0, \end{cases} \end{aligned} \quad (4.9)$$

$$\begin{aligned}
& \mathcal{A}_k^2(x_1, x_2, T_n) \\
&= \sum_{j=n}^N \frac{c_j \bar{A}(T_n, T_{j+1})}{(u_k)^2 + (\bar{B}(T_n, T_{j+1}))^2} \\
&\quad \times [\exp\{-\bar{B}(T_n, T_{j+1})x_2\} \\
&\quad \times (u_k \sin(u_k(x_2 - a)) - \bar{B}(T_n, T_{j+1}) \cos(u_k(x_2 - a))) \\
&\quad - \exp\{-\bar{B}(T_n, T_{j+1})x_1\} \\
&\quad \times (u_k \sin(u_k(x_1 - a)) - \bar{B}(T_n, T_{j+1}) \cos(u_k(x_1 - a)))]. \quad (4.10)
\end{aligned}$$

Here, $c_j = \tau_j K$, $j = n, \dots, N-1$, $c_N = 1 + \tau_N K$, and \bar{A} and \bar{B} are coefficients associated to the ZCB price, given in Appendix 1 (available online).

In the COS method, computation also takes place in backward fashion. We distinguish an early exercise date from an intermediate date between two exercise times. At an intermediate date, $t_m \in (T_{n-1}, T_n)$, $V(t_m, \cdot) = C(t_m, \cdot)$; thus, the coefficients \mathcal{V}_k at time t_m are given by

$$\mathcal{V}_k(t_m) = \int_a^b C(t_m, y) \cos(u_k(y - a)) dy = \mathcal{C}_k(a, b, t_m), \quad (4.11)$$

where the coefficients \mathcal{C}_k at time t_m over $[x_1, x_2]$ are computed via an integral

$$\begin{aligned}
\mathcal{C}_k(x_1, x_2, t_m) &:= \int_{x_1}^{x_2} C(t_m, y) \cos(u_k(y - a)) dy \\
&\approx \frac{2}{b-a} \sum_{j=0}^{Q-1} \left(\int_{x_1}^{x_2} \mathcal{P}_j(y, t_m, t_{m+1}) \cos(u_k(y - a)) dy \right) \mathcal{V}_j(t_{m+1}) \\
&= \frac{2}{b-a} \sum_{j=0}^{Q-1} \operatorname{Re}\{\mathcal{W}_j(t_m, t_{m+1}) \mathcal{X}_{kj}(x_1, x_2, \Delta t_m)\} \mathcal{V}_j(t_{m+1}),
\end{aligned} \quad (4.12)$$

in which coefficients $\{\mathcal{V}_j(t_{m+1})\}_{j=1}^{Q-1}$ have been computed at time t_{m+1} , and coefficients \mathcal{W} and \mathcal{X} are given by

$$\begin{aligned}
& \mathcal{W}_j(t_m, t_{m+1}) \\
&= \exp \left\{ - \int_{t_m}^{t_{m+1}} \theta(s) ds + iu_j \theta(t_{m+1}) \right. \\
&\quad \left. - iu_j a - \tilde{B}_g(u_j, \Delta t_m) \theta(t_m) + \tilde{A}_g(u_j, \Delta t_m) \right\}, \quad (4.13)
\end{aligned}$$

$$\begin{aligned}
& \mathcal{X}_{kj}(x_1, x_2, \Delta t_m) \\
&= \int_{x_1}^{x_2} \cos(u_k(y-a)) \exp(\tilde{B}_g(u_j, \Delta t_m)y) dy \\
&= \frac{1}{(u_k)^2 + (\tilde{B}_g(u_j, \Delta t_m))^2} \\
&\quad \times [\exp\{\tilde{B}_g(u_j, \Delta t_m)x_2\} \\
&\quad \times (u_k \sin(u_k(x_2-a)) + \cos(u_k(x_2-a))\tilde{B}_g(u_j, \Delta t_m)) \\
&\quad - \exp\{\tilde{B}_g(u_j, \Delta t_m)x_1\} \\
&\quad \times (u_k \sin(u_k(x_1-a)) + \cos(u_k(x_1-a))\tilde{B}_g(u_j, \Delta t_m))]. \quad (4.14)
\end{aligned}$$

The analytic formulas for the coefficients \tilde{A}_g and \tilde{B}_g , and the integral with the function θ in (4.13), are given in Appendix 1 (available online).

At an early exercise date $t_m = T_n$, $n = N - 1, \dots, 1$, the option value is the maximum of the continuation value and the immediate exercise value; hence, we solve the following equation: $C(T_n, x^*(T_n)) - U_n(x^*(T_n)) = 0$. Solution $x^*(T_n)$ represents the optimal early exercise boundary at time T_n . The equation can be solved by some root-finding algorithm, such as the Newton–Raphson method.

The coefficients $\{\mathcal{V}_k(t_m)\}_{k=0}^{Q-1}$ at time $t_m = T_n$ with the optimal exercise value $x^*(T_n)$ are given by

$$\begin{aligned}
\mathcal{V}_k(T_n) &= \int_a^b \max(C(T_n, y), U_n(y)) \cos(u_k(y-a)) dy \\
&= \begin{cases} \mathcal{G}_k(r^*(T_n), b, T_n) + \mathcal{C}_k(a, r^*(T_n), T_n) & \text{payer,} \\ \mathcal{G}_k(a, r^*(T_n), T_n) + \mathcal{C}_k(r^*(T_n), b, T_n) & \text{receiver.} \end{cases} \quad (4.15)
\end{aligned}$$

The computation of the coefficients $\{\mathcal{V}_k\}_{k=0}^{Q-1}$ depends on the early exercise boundary value at each exercise date T_n . The continuation function from the Fourier-cosine expansions in (4.5) converges with respect to the number of Fourier terms Q when the integration interval is chosen properly.

At each t_m , the continuation values for all scenarios can be computed by (4.5). Risk-neutral and real-world exposure profiles are obtained by backward iteration, as in Sections 3.1.1 and 3.1.2. One can employ interpolation to enhance the computational speed for the computation of the continuation values.

5 NUMERICAL EXPERIMENTS

We test the developed algorithms for different test cases under the one-factor Hull–White and two-factor G2++ models.

The notional amount of the underlying swap is set equal to 100. We define a $T_1 \times T_N$ Bermudan swaption as a Bermudan option written on the underlying swap that can be exercised between T_1 and T_N , the first and last early exercise opportunities. The option can be exercised annually after T_1 , and the payment of the underlying swap is made by the end of each year until the fixed end date $T_{N+1} = T_N + 1$, ie, the time fraction $\tau_n = 1$.

We take a fixed strike K to be $40\%S$, $100\%S$ and $160\%S$, where S is the swap rate associated with date T_1 and payment dates $\mathcal{T}_1 = \{T_2, \dots, T_{N+1}\}$ given by (2.9). It is the at-the-money strike of the European swaption that expires at date T_1 associated with payment dates \mathcal{T}_1 .

5.1 Experiments with the Hull–White model

We generate risk-neutral and real-world scenarios using the Hull–White model presented in Appendix 1 (available online), with risk-neutral parameters λ and η obtained by market prices and real-world parameters κ and η obtained by historical data.

Table 1 reports the time-zero option values, CVA and real-world EPEs and MPFEs of 1Y×5Y and 4Y×10Y receiver Bermudan swaptions by the COS method, SGBM and LSM-bundle and LSM-all algorithms.

For the computation of future exposure distributions, one needs to combine the COS method computations with Monte Carlo scenario generation, so there are standard errors as well for the corresponding CVA, EPE and MPFE values. We present $100 \times$ CVA values instead of CVA to enlarge the differences and standard errors in Table 1.

The reference results by the COS method are obtained with $Q = 100$ cosine terms. In the SGBM algorithm, we use as basis functions $\{1, r, r^2\}$ for the approximation of the continuation values and ten bundles containing an equal number of paths. In the LSM, we choose a cubic function based on $\{1, r, r^2, r^3\}$ for the approximation. It is observed that SGBM and LSM converge with respect to the number of basis functions, and, from our experiments, we also find that for longer maturities a larger number of basis functions are required to maintain the accuracy.

As shown in Table 1, the differences in the computed time-zero option values between all algorithms are very small. The LSM-bundle and LSM-all algorithms return the same time-zero option value, as they are based on the same technique to determine the early exercise policy. Compared with the LSM, the SGBM has improved accuracy with smaller variances. The absolute difference in V_0 -values between the SGBM and COS method is as small as 10^{-3} , and the standard errors are less than 1%. The largest difference in V_0 between the LSM and COS method is 6×10^{-3} , with a standard error between 1 and 2% in Table 1.

The SGBM is particularly accurate for computing the MPFE values. The results in Table 1 show that the absolute differences for MPFE computed by the SGBM and COS methods are less than 0.01. The LSM-all algorithm does not result in satisfactory results for the exposure values. MPFE is overestimated, while EPE is underestimated. The LSM-bundle algorithm, however, shows significant improvements with smaller errors.

For the computation of EPE and MPFE, the results obtained via these algorithms have a similar standard error. This shows that the dominating factor in the EPE and MPFE variances is connected to the number of generated scenarios.

Figure 3 compares the statistics of the risk-neutral and real-world exposure distributions: the mean in Figure 3(a) and the 99% quantile in Figure 3(b) for a 4Y/10Y receiver Bermudan swaption along time horizon $[0, 10]$. The significant difference between the curves shows that one cannot use quantiles computed by risk-neutral exposure distributions to represent the real-world PFE. There are downward jumps in the EE and PFE curves at each early exercise date $\{4Y, 5Y, \dots\}$ as the swaption on some of the paths is exercised.

The mean and 99% percentile of the real-world exposure are the required EE and PFE values. Figure 4 compares the EE and PFE curves obtained by the different algorithms for the periods 2Y–4Y and 6Y–8Y. The LSM tends to overestimate the PFE prior to the first early exercise opportunity and underestimate it afterwards. The SGBM results are as accurate as the reference values.

The main reason for the SGBM's excellent fit in the tails of the distributions is that, at each date, the algorithm provides an accurate local approximation of the continuation function for the whole realized domain of the underlying factor.

Figure 5(a) compares the reference continuation functions (by COS) with the approximated continuation functions (by the SGBM and LSM) on the bounded realized risk-neutral region at time 6.5Y, for the 4Y/10Y receiver swaption. The approximation by LSM-all is not accurate at the upper and lower regions, which explains its performance in Table 1. We observe an accuracy improvement in the results from the LSM-bundle algorithm. From the plot, we observe that the SGBM's approximated function well resembles the reference value on the whole domain. Figure 5(b) presents the empirical density of the risk-neutral short rate and the observed real-world short rate, where we see that the realized domain under the risk-neutral measure is more widely spread.

Table 2 gives the computational times for these algorithms. The SGBM is significantly faster than the reference COS algorithm, while the LSM is less accurate but faster than the SGBM. The experiments are performed on a computer with a CPU Intel Core i7-2600 3.40GHz \times 8 processor and 15.6 Gigabytes of RAM. The computational cost increases proportionally with respect to parameter M .

TABLE 1 Bermudan receiver swaption under the Hull–White model.

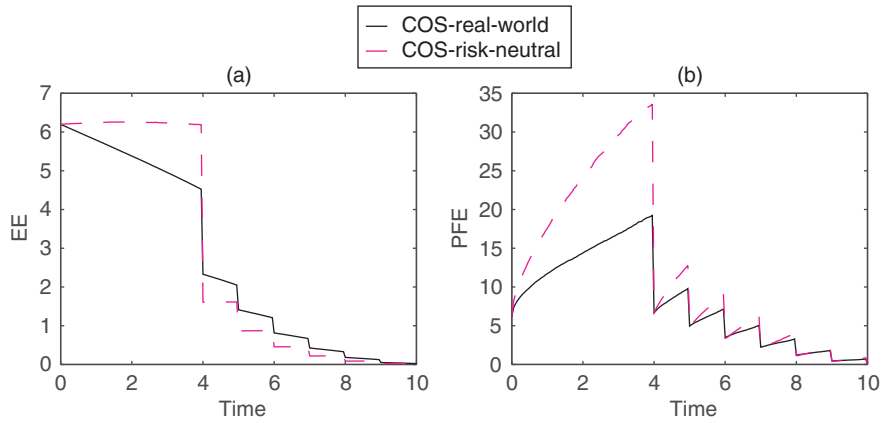
(a) 1Y×5Y					
<i>K/S</i>	Value	COS	SGBM	LSM-bundle	LSM-all
40%	V_0	4.126	4.127(0.00)	4.126(0.01)	4.126(0.01)
	MPFE	9.125(0.06)	9.118(0.06)	9.039(0.05)	8.8(0.05)
	EPE	1.704(0.00)	1.705(0.00)	1.708(0.01)	1.806(0.01)
	100CVA	15.87(0.01)	15.74(0.01)	15.75(0.08)	15.87(0.08)
100%	V_0	5.463	5.464(0.00)	5.461(0.01)	5.461(0.01)
	MPFE	11.07(0.05)	11.07(0.05)	11.06(0.05)	10.9(0.04)
	EPE	2.094(0.00)	2.096(0.00)	2.098(0.00)	2.215(0.00)
	100CVA	18.56(0.02)	18.33(0.02)	18.35(0.04)	18.44(0.04)
160%	V_0	7.11	7.11(0.00)	7.113(0.01)	7.113(0.01)
	MPFE	14.43(0.04)	14.42(0.04)	14.26(0.04)	13.92(0.04)
	EPE	2.368(0.00)	2.369(0.00)	2.372(0.01)	2.483(0.01)
	100CVA	21.28(0.02)	20.95(0.02)	20.98(0.05)	21.01(0.05)
(b) 4Y×10Y					
<i>K/S</i>	Value	COS	SGBM	LSM-bundle	LSM-all
40%	V_0	4.235	4.236(0.00)	4.237(0.01)	4.237(0.01)
	MPFE	14.12(0.12)	14.13(0.12)	13.86(0.11)	13.16(0.09)
	EPE	1.827(0.00)	1.829(0.00)	1.834(0.01)	1.91(0.01)
	100CVA	38.22(0.02)	37.98(0.02)	38.04(0.13)	38.34(0.13)
100%	V_0	6.199	6.199(0.00)	6.201(0.02)	6.201(0.02)
	MPFE	19.29(0.11)	19.29(0.11)	19.08(0.12)	18.08(0.10)
	EPE	2.606(0.00)	2.607(0.00)	2.616(0.01)	2.719(0.01)
	100CVA	53.35(0.05)	52.92(0.05)	53.03(0.14)	53.26(0.14)
160%	V_0	8.691	8.691(0.00)	8.687(0.02)	8.687(0.02)
	MPFE	24.33(0.09)	24.34(0.09)	24.28(0.09)	23.42(0.09)
	EPE	3.526(0.00)	3.527(0.00)	3.539(0.01)	3.628(0.01)
	100CVA	71.94(0.06)	71.35(0.06)	71.47(0.09)	71.5(0.10)

(a) $S \approx 0.0109$; risk-neutral: $\sigma = 0.010$, $\lambda = 0.020$; real-world: $\eta = 0.010$, $\kappa = 0.015$. (b) $S \approx 0.0113$; risk-neutral: $\sigma = 0.020$, $\lambda = 0.012$; real-world: $\eta = 0.006$, $\kappa = 0.008$. Risk-neutral and real-world scenarios are generated; the forward rate is flat, $f^M(0, t) = 0.01$; the default probability function $PS(t) = 1 - \exp(-0.02t)$ and $LGD = 1$; option values and CVA are based on $K_d = 100 \times 10^3$ risk-neutral scenarios; MPFE and EPE are based on $K_a = 100 \times 10^3$ real-world scenarios; the number of monitoring dates $M = T_N / \Delta t$ with $\Delta t = 0.05$; standard errors are in parentheses, based on ten independent runs.

5.2 Experiments with the G2++ model

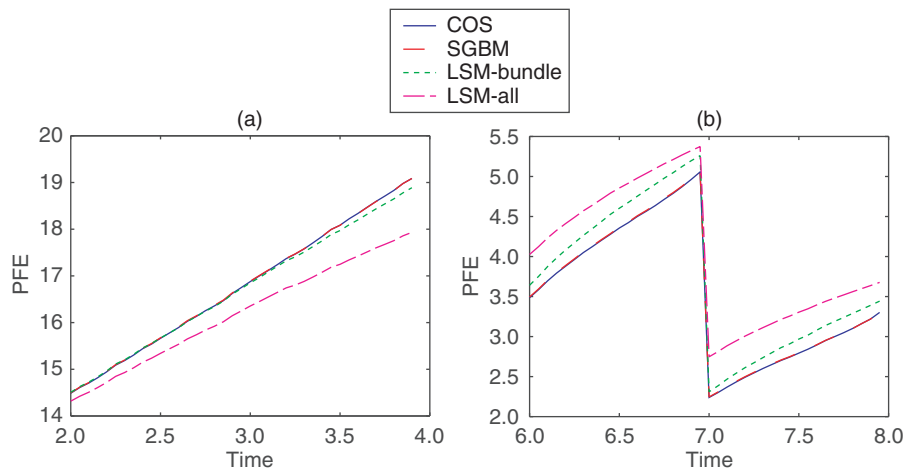
The dynamics of the risk-neutral and real-world G2++ models are given in Appendix 2 (available online), where the associated parameters (ie, the reversion speed κ_1 , κ_2 ; the

FIGURE 3 Comparison of the mean and 99% quantile of the exposure distributions, computed by the COS method and based on risk-neutral and real-world scenarios, for the 4Y/10Y Bermudan receiver swaption, as specified in Table 1, when $K/S = 1$.



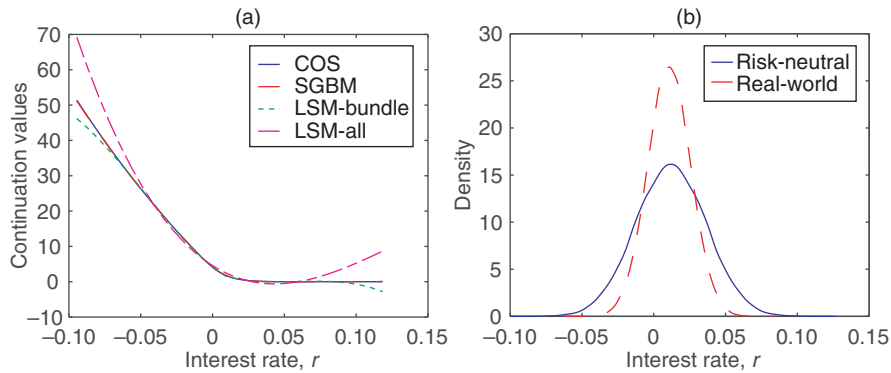
(a) Exposure average. (b) Exposure quantile 99%.

FIGURE 4 Comparison of PFE curves obtained by the COS method, SGBM and LSM for 2Y–4Y and 6Y–8Y, for the 4Y/10Y Bermudan receiver swaption specified in Table 1, when $K/S = 1$.



(a) PFE, 2Y–4Y. (b) PFE, 6Y–8Y.

FIGURE 5 Comparison of continuation functions via all algorithms at 6.5Y, for the 4Y/10Y Bermudan receiver swaption, specified in Table 1, when $K/S = 1$.



(a) Continuation function. (b) Short rate density.

TABLE 2 Computational costs (seconds) for computation of risk-neutral and real-world exposure distributions.

	T_N	M	COS (s)	SGBM (s)	LSM-bundle (s)	LSM-all (s)
	5Y	100	60.9	4.48	3.32	1.67
	10Y	200	122.7	9.02	7.09	3.26

volatility η_1, η_2 ; and the correlation ξ) are based on historical data, and risk-neutral parameters ($\lambda_1, \lambda_2, \sigma_1, \sigma_2$ and ρ) are based on market prices.

In this two-dimensional model, we use the following monomials as the basis functions in the LSM algorithm

$$\{1, x_t, y_t, x_t^2, x_t y_t, y_t^2, x_t^3, x_t^2 y_t, x_t y_t^2, y_t^3\};$$

the basis functions in the SGBM algorithm are given by

$$\{1, x_t, y_t, x_t^2, x_t y_t, y_t^2\},$$

from which we observe that the number of basis functions increases rapidly with regard to the dimension of the underlying variable.

The associated discounted moments, required in the SGBM, can easily be derived from the analytic formula of the dChF of the G2++ model. As for the Hull–White model, we use $J = 10$ bundles in SGBM. In SGBM, we can either use the

TABLE 3 Receiver Bermudan swaption under the G2++ model.

(a) 1Y×5Y				
<i>K/S</i>	Value	SGBM	LSM-bundle	Difference
40%	V_0	1.742(0.00)	1.747(0.01)	0.005
	MPFE	5.066(0.02)	5.037(0.17)	-0.029
	EPE	0.771(0.00)	0.773(0.00)	0.002
	100CVA	7.491(0.01)	7.51(0.03)	0.019
100%	V_0	2.897(0.00)	2.900(0.01)	0.003
	MPFE	6.535(0.04)	6.448(0.11)	-0.087
	EPE	1.113(0.00)	1.113(0.01)	0.000
	100CVA	10.05(0.01)	10.07(0.03)	0.02
160%	V_0	4.560(0.00)	4.563(0.01)	0.003
	MPFE	9.652(0.01)	9.601(0.09)	-0.51
	EPE	1.33(0.00)	1.337(0.01)	0.007
	100CVA	12.67(0.01)	12.7(0.04)	0.03

(b) 3Y×10Y				
<i>K/S</i>	Value	SGBM	LSM-bundle	Difference
40%	V_0	0.861(0.00)	0.865(0.00)	0.005
	MPFE	2.784(0.03)	2.704(0.04)	-0.08
	EPE	0.446(0.00)	0.447(0.00)	0.001
	100CVA	8.678(0.01)	8.705(0.03)	0.027
100%	V_0	2.466(0.00)	2.475(0.01)	0.008
	MPFE	7.176(0.03)	7.115(0.03)	-0.061
	EPE	1.059(0.00)	1.063(0.00)	0.004
	100CVA	19.53(0.01)	19.61(0.04)	0.08
160%	V_0	5.42(0.00)	5.428(0.00)	0.008
	MPFE	12.1(0.03)	12.22(0.04)	0.012
	EPE	1.839(0.00)	1.846(0.00)	0.007
	100CVA	35.49(0.01)	35.62(0.04)	0.13

(a) $S \approx 0.0104$; risk-neutral: $\sigma_1 = 0.015$, $\sigma_2 = 0.008$, $\lambda_1 = 0.07$, $\lambda_2 = 0.08$, $\rho = -0.6$; real-world: $\eta_1 = 0.005$, $\eta_2 = 0.01$, $\kappa_1 = 0.54$, $\kappa_2 = 0.07$, $\xi = -0.8$. (b) $S \approx 0.0102$; risk-neutral: $\sigma_1 = 0.005$, $\sigma_2 = 0.008$, $\lambda_1 = 0.09$, $\lambda_2 = 0.15$, $\rho = -0.6$; real-world: $\eta_1 = 0.002$, $\eta_2 = 0.006$, $\kappa_1 = 0.04$, $\kappa_2 = 0.07$, $\xi = -0.8$. Risk-neutral and real-world scenarios are generated; forward rate $f^M(0, t) = 0.01$; the default probability function $PS(t) = 1 - \exp(-0.02t)$ and $LGD = 1$; option values and CVA are based on $K_q = 100 \times 10^3$ risk-neutral scenarios; MPFE and EPE are based on $K_q = 100 \times 10^3$ real-world scenarios; the number of monitoring dates $M = T_N / \Delta t$ with $\Delta t = 0.05$.

two-dimensional equal-number bundling method, introduced in Feng and Oosterlee (2014), or the one-dimensional version based on projecting the high-dimensional variable onto a one-dimensional variable. Here, we create the bundles based on the realized values of $(x_t + y_t)$ on each path at time t_m .

FIGURE 6 PFE and the 99% quantile of the exposure distributions of a receiver Bermudan swaption, as specified in Table 3, when $K/S = 1$.

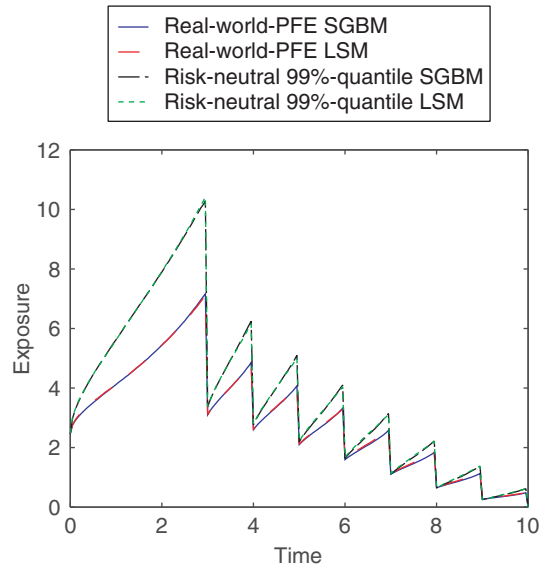


TABLE 4 Computational costs (seconds) for the computation of risk-neutral and real-world exposure distributions under the G2++ model.

T_N	M	LSM-bundle (s)	SGBM (s)
5Y	100	6.28	8.10
10Y	200	12.96	15.07

Table 3 reports the time-zero option value results for SGBM and LSM as well as the exposure measures for receiver Bermudan swaptions, where we can analyze the difference between the results by these algorithms.

Figure 6 presents the PFE curves computed on the real-world scenarios as well as the mean and 99% quantiles of the risk-neutral exposure distributions at each monitoring date. As expected, there is a clear difference between the statistics of the risk-neutral and real-world exposure distributions.

Table 4 presents the computational cost of the algorithms for this two-dimensional model. The cost increases with respect to the dimension of the variable.

6 CONCLUSION

This paper presents computationally efficient techniques for the simultaneous computation of exposure distributions under the risk-neutral and observed real-world probability measures. They are based on only two sets of scenarios, one generated under the risk-neutral dynamics and another under the observed real-world dynamics, as well as on basic techniques such as regression. Compared with nested Monte Carlo simulation, the techniques presented significantly reduce the computational cost and maintain high accuracy, which we demonstrated by using numerical results for Bermudan swaptions and comparing these with reference results generated by the Fourier-based COS method. We illustrated the ease of implementation for both the one-factor Hull–White and two-factor G2++ models.

We recommend the SGBM because of its accuracy and efficiency in the computation of continuation values. A highly satisfactory alternative is to use the LSM-bundle approach. The reference COS method is highly efficient for computing time-zero values of the Bermudan swaption, but for the computation of exposure, there is room for improvement in terms of computational speed.

The results for the parameter values chosen show that there are clear differences in exposure distributions for the risk-neutral and real-world scenarios. The proposed algorithms are based on the requirement that the sample space induced by the observed historical model is a subspace of the sample space under the risk-neutral measure.

The valuation framework presented is flexible and may be used efficiently for any type of Bermudan-style claim, such as Bermudan options and swaptions. For a Bermudan option, one can compute the sensitivities of CVA at the same time as using the SGBM, which is an additional benefit. The algorithms developed can be extended easily to the situation in which model parameters are piecewise constant over the time horizon.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the contents of the paper. The views expressed in this paper are those of the authors and do not necessarily reflect the position of their employers. Financial support from the Dutch Technology Foundation STW (project 12214) is gratefully acknowledged.

REFERENCES

- Andersen, L. B. (1999). A simple approach to the pricing of Bermudan swaptions in the multifactor Libor market model. *The Journal of Computational Finance* 3(2), 5–32 (<http://doi.org/bkn9>).

- Andersen, L. B., and Piterbarg, V. V. (2010). *Interest Rate Modeling*. Atlantic Financial Press.
- Basel Committee on Banking Supervision (2005). Annex 4 to “International convergence of capital measurement and capital standards: a revised framework”. Report, Bank for International Settlements.
- Basel Committee on Banking Supervision (2010). Basel III: a global regulatory framework for more resilient banks and banking systems. Report, Bank for International Settlements.
- Brigo, D., and Mercurio, F. (2007). *Interest Rate Models – Theory and Practice: With Smile, Inflation and Credit*. Springer Science & Business Media.
- Fang, F., and Oosterlee, C. W. (2009). Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions. *Numerische Mathematik* **114**(1), 27–62 (<http://doi.org/b2vhnX>).
- Feng, Q., and Oosterlee, C. W. (2014). Monte Carlo calculation of exposure profiles and Greeks for Bermudan and barrier options under the Heston Hull–White Model. SSRN Working Paper, arXiv:1412.3623.
- Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*, Volume 53. Springer Science & Business Media (<http://doi.org/bf7w>).
- Gregory, J. (2010). *Counterparty Credit Risk: The New Challenge for Global Financial Markets*, Volume 470. Wiley.
- Hull, J. C., Sokol, A., and White, A. (2014). Modeling the short rate: the real and risk-neutral worlds. Working Paper 2403067, Rotman School of Management (<http://doi.org/bkpb>).
- Jain, S., and Oosterlee, C. W. (2012). Pricing high-dimensional Bermudan options using the stochastic grid method. *International Journal of Computer Mathematics* **89**(9), 1186–1211 (<http://doi.org/bkpc>).
- Jain, S., and Oosterlee, C. W. (2015). The stochastic grid bundling method: efficient pricing of Bermudan options and their Greeks. *Applied Mathematics and Computation* **269**, 412–431 (<http://doi.org/bkpd>).
- Joshi, M. S., and Kwon, O. K. (2016). Least squares Monte Carlo credit value adjustment with small and unidirectional bias. SSRN Working Paper 2717250 (<http://doi.org/bkpf>).
- Karlsson, P., Jain, S., and Oosterlee, C. W. (2014). Counterparty credit exposures for interest rate derivatives using the stochastic grid bundling method. SSRN Working Paper 2538173 (<http://doi.org/bkpg>).
- Kenyon, C., Green, A. D., and Berrahoui, M. (2015). Which measure for PFE? The risk appetite measure A. SSRN Working Paper, December 15, arXiv:1512.06247.
- Leitao, Á., and Oosterlee, C. W. (2015). GPU acceleration of the stochastic grid bundling method for early-exercise options. *International Journal of Computer Mathematics* **92**(12), 2433–2454 (<http://doi.org/bkph>).
- Longstaff, F. A., and Schwartz, E. S. (2001). Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies* **14**(1), 113–147 (<http://doi.org/b38b5q>).
- Øksendal, B. (2003). *Stochastic Differential Equations*. Springer (<http://doi.org/dqpdqb>).
- Ruijter, M. J., and Oosterlee, C. W. (2012). Two-dimensional Fourier cosine series expansion method for pricing financial options. *SIAM Journal on Scientific Computing* **34**(5), B642–B671 (<http://doi.org/bkpj>).
- Ruiz, I. (2012). Backtesting counterparty risk: how good is your model? Technical Report, iRuiz Consulting.

Stein, H. J. (2013). Joining risks and rewards. SSRN Working Paper 2368905.

Stein, H. J. (2014). Fixing underexposed snapshots: proper computation of credit exposures under the real world and risk neutral measures. SSRN Working Paper 2365540 (<http://doi.org/bkpk>).

Zhu, S. H., and Pykhtin, M. (2007). A guide to modeling counterparty credit risk. *GARP Risk Review* July/August(37), 16–22.