



A static replication approach for callable interest rate derivatives: mathematical foundations and efficient estimation of SIMM–MVA

J. H. Hoencamp, S. Jain & B. D. Kandhai

To cite this article: J. H. Hoencamp, S. Jain & B. D. Kandhai (26 Feb 2024): A static replication approach for callable interest rate derivatives: mathematical foundations and efficient estimation of SIMM–MVA, *Quantitative Finance*, DOI: [10.1080/14697688.2024.2312523](https://doi.org/10.1080/14697688.2024.2312523)

To link to this article: <https://doi.org/10.1080/14697688.2024.2312523>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 26 Feb 2024.



[Submit your article to this journal](#)



Article views: 312



[View related articles](#)



[View Crossmark data](#)

A static replication approach for callable interest rate derivatives: mathematical foundations and efficient estimation of SIMM–MVA

J. H. HOENCAMP*†§, S. JAIN‡ and B. D. KANDHAI†§

†Computational Science Lab, University of Amsterdam, Science Park 904, Amsterdam 1098XH, Netherlands

‡Department of Management Studies, Indian Institute of Science, Bangalore, India

§Quantitative Analytics, ING Bank, Foppingadreef 7, Amsterdam, 1102 BD, Netherlands

(Received 16 June 2023; accepted 20 January 2024; published online 26 February 2024)

The computation of credit risk measures such as exposure and Credit Value Adjustments (CVA) requires the simulation of future portfolio prices. Recent metrics, such as dynamic Initial Margin (IM) and Margin Value Adjustments (MVA) additionally require the simulation of future conditional sensitivities. For portfolios with non-linear instruments that do not admit closed-form valuation formulas, this poses a significant computational challenge. This problem is addressed by proposing a static replication algorithm for interest rate options with early-exercise features under an affine term-structure model. Under the appropriate conditions, we can find an equivalent portfolio of vanilla options that replicate these products. Specifically, we decompose the product into a portfolio of European swaptions. The weights and strikes of the portfolio are obtained by regressing the target option value with interpretable, feed-forward neural networks. Once an equivalent portfolio of European swaptions is determined, we can leverage on closed-form expressions to obtain the conditional prices and sensitivities, which serve as an input to exposure and SIMM-driven MVA quantification. For a consistent forward sensitivity estimation, this involves the differentiation of the portfolio-weights. The accuracy and convergence of the method is demonstrated through several representative numerical examples, benchmarked against the established least-square Monte Carlo method.

Keywords: Static replication; Bermudan swaptions; Affine term-structure modeling; Counterparty credit risk; Initial Margin; CVA; MVA

1. Introduction

This research contributes to the quantification of risk measures related to *counterparty credit risk* (CCR) for interest rate derivatives that are traded over-the-counter (OTC). CCR refers to the risk that the counterparty in a bilateral agreement may default before the final settlement of every cash-flow and will fail to meet its financial obligations (Bank for International Settlements 2019). This type of risk has gained a lot of attention in the aftermath of the global financial crisis of 2007–08. Since then, the regulators have introduced numerous measures to mitigate systemic credit risk that have resulted in quantifying CCR through various value adjustments (xVA) (Gregory 2020). The family of xVAs includes CVA, which is defined as the expected loss that is incurred from a counterparty defaulting. Other xVAs quantify the

cost related to funding, collateral and capital requirements. A general introduction to the challenges related to CCR quantification can be found in Zhu and Pykhtin (2007).

An important means to reduce CCR for uncleared OTC derivatives is to post collateral (Gregory 2020). For interest rate derivatives we can in general distinguish two types of collateral: *Variation Margin* (VM) and *Initial Margin* (IM). VM matches the current value of the underlying portfolio and needs to be updated on a regular basis. The funding cost related to VM is known as FVA and is reflected by the expected exposure of the portfolio (Burgard and Kjaer 2013), similar to CVA. Posting VM significantly reduces the exposure to counterparty risk but does not bring it down to zero. Typically there will be a delay between the event of default and the settlement of all outstanding positions, which is known as the margin period of risk (Gregory 2020). IM serves as a protection against exposure changes due to market moves during this close-out period. In general, IM is supposed to

*Corresponding author. Email: j.h.hoencamp@uva.nl

cover the potential future exposure of a contract over a time-horizon of 10 business days with a 99% confidence. In that regard, the volume of IM typically reflects the 99% value-at-risk of the 10-day move in portfolio value. A portfolio's VaR varies over time and should therefore (just like VM) be updated on a regular basis.

A VaR measure is highly model-dependent and computationally intensive to quantify. To address these problems, the International Swaps and Derivatives Association (ISDA) has developed the *Standard Initial Margin Model* (SIMM). Since its publication in 2016, SIMM has become an industry-standard approach to IM quantification. Benefits of such a common methodology include a transparent dispute resolution and a consistent regulatory governance of collateral (International Swaps and Derivatives Association, Inc 2016). In short, SIMM is a sensitivity-based approach, in which VaR measures are approximated using the instrument's sensitivity to the shock of an underlying risk-factor. For the user of SIMM, this means that IM quantification reduces to the calculation of a set of bucketed portfolio sensitivities, such as Delta and Vega.

The exchange of collateral needs to be funded, which comes at a cost for the dealer. The total expected funding cost over the life-time of a derivative contract is known as the *Margin Value Adjustment* (MVA) (Green and Kenyon 2015) and is a recent addition to the xVA-collection. An ISDA margin survey reported that market participants collected over \$300 billion of IM for non-cleared derivatives at year-end 2021 (International Swaps and Derivatives Association, Inc 2021). Due to its volume, an adequate quantification and management of IM-induced funding cost is highly relevant. Under the assumption that the size of IM is implied by SIMM and thus by sensitivities, it follows that MVA is driven by the future distribution of portfolio sensitivities over the life-time of the trade. Within a Monte Carlo simulation framework, this means sensitivities need to be calculated along the Monte Carlo path. For vanilla derivatives, which admit closed-form Greeks, the computational cost may be manageable. However, for exotic derivatives, such as options with early-exercise features, this is a demanding problem.

In this work, we present a universal approach to the computation of prices, Greeks, exposures and sensitivities along the Monte Carlo path for callable interest rate derivatives. Our focus will be on Bermudan swaptions, which is a class of exotic OTC derivatives that is heavily traded in the market. Our methodology is an extension of the work presented in Lokeshwar *et al.* (2022) and Hoencamp *et al.* (2023) and relies on the concept of *static replication*. A static replication is a portfolio of vanilla instruments, that mirrors the value of the original exotic option, until it is either exercised or matured. The portfolio composition of a static replication is constant throughout the life-time of the trade. This is in contrast with a dynamic replication, which needs to be continuously re-balanced (such as a traditional Delta hedge), or a semi-static replication, which needs to be re-balanced on a finite number of instances. The decomposition of a complex product into a portfolio of vanilla options greatly simplifies its risk analysis and allows for an efficient price and sensitivity calculation.

In the context of CVA, methodologies that facilitate efficient exposure calculations, have received a lot of

attention in the recent literature. Ordinary least-square regression techniques for American-style options embedded in a simulation framework, have been popularized by Carriere (1996), Tsitsiklis and Van Roy (2001) and Longstaff and Schwartz (2001). For exposure purposes, this regression methodology is extended by Joshi and Kwon (2016) and Feng *et al.* (2016) to reduce regression bias and noise along the Monte Carlo path. In other works, Monte Carlo methods have been combined with finite-difference estimation (De Graaf *et al.* 2014, Simaitis *et al.* 2016), the COS-method (Shen *et al.* 2013, Feng *et al.* 2016), the Stochastic Grid Bundling Method (SGBM) (Karlsson *et al.* 2016) and Chebyshev interpolation (Glau *et al.* 2021). In recent advances, machine learning has been applied to facilitate efficiency in exposure calculations. This includes a Deep Optimal Stopping algorithm to estimate the exercise strategy (Andersson and Oosterlee 2021) and a Deep xVA Solver, which relies on a neural network-based BSDE solver.

The publication of SIMM and the increasing significance of MVA have been incentives to investigate methodologies that facilitate efficient sensitivity calculation. For practitioners, the default approach is 'bump-and-reval', due to its simplicity and straight-forward implementation. For spot calculations, this is feasible, but for an application such as MVA, it would imply path-wise model re-calibration and option re-evaluation for each scenario, time-step and sensitivity-component. Thus, the computational burden would be excessive. In Glasserman (2004), unbiased Monte Carlo sensitivity estimators are proposed, which yield better convergence properties than the traditional finite-difference estimator. These are known as the *likelihood ratio method*, which relies on differentiating the payoff probability density, and the *path-wise derivative method*, which relies on differentiating the payoff function itself. Yet the computational improvement of these methods can be limited for path-dependent instruments.

The application of *algorithmic differentiation* (AD) to the path-wise derivative method has been shown to yield a significant efficiency gain. The concept of AD is that the estimation routine is decomposed into a series of basic operations, which are differentiated by a repeated application of the chain rule (Griewank and Walther 2008). The differentiation is either performed in a forward iteration (tangent mode) or backward iteration (adjoint mode, AAD). See for example Giles and Glasserman (2006), Capriotti (2011) or Capriotti *et al.* (2017) for an application with American-style derivatives. The tangent mode is efficient if the number of independent variables (i.e. the dimensionality of the function input) is small. It is therefore less suitable for the calculations of large gradients. In contrast, the adjoint mode is efficient if the number of dependent variables (i.e. the dimensionality of the function output) is small. A limitation of AAD, however, is that it is demanding in terms of memory and its implementation is often technical. Others have combined adjoint differentiation with LSMC (Antonov *et al.* 2017, Caspers and Lichters 2018), stochastic automatic differentiation (Fries 2019) or SGBM (Jain *et al.* 2019). Alternative approaches to path-wise sensitivity approximations include principal component analysis (Kappen 2017), Chebyshev interpolation (Zeron and Ruiz 2018) and least-square regression (Lakhany and Zhang 2021).

With a focus on Bermudan swaptions under an affine term-structure model, our contribution to the literature is three-fold. First we extend the semi-static replication algorithm that is presented in Lokeshwar *et al.* (2022) and Hoencamp *et al.* (2023) to a replication approach that is truly static. The method is based on formulating the portfolio optimization as the regression of a shallow neural network to the target’s option value. As a result, the trained parameters represent the weights and strikes of the portfolio composition. Second, we show that with the proper constraints on the regression, a Bermudan can be replicated with a portfolio of European swaptions. The accuracy is demonstrated through several numerical examples. Thirdly, we derive efficient estimators for sensitivities along the Monte Carlo path, which serve as an input to dynamic IM quantification. Here we exploit the fact that Deltas and Vegas of European swaptions are available in closed-form. We also show that for complete and accurate sensitivities, a differentiation of the portfolio weights needs to be incorporated. We demonstrate the performance of the algorithm through several numerical examples and convergence analyzes, benchmarking the results to the established least-square Monte Carlo method (Longstaff and Schwartz 2001).

The outline of the article is as follows. In section 2, we define the modeling landscape, the considered interest rate

derivatives and the risk metrics of interest. In section 3, we present the regression algorithm and show how we can achieve a static replication for Bermudan swaptions. In section 4, we derive the estimators used for price and sensitivity calculation along the Monte Carlo path. Section 5 is subject to several representative numerical examples. We conclude with a summary of the results in section 6.

2. Problem formulation

A problem overview is provided in figures 1–4.

2.1. Problem overview

- (1) Initial Margin (IM) is a form of collateral, which serves as a protection against the 99% value-at-risk of the exposure change during the margin period of risk (figure 1).
- (2) The industry standard approach to IM quantification is the Standard Initial Margin Model (SIMM), which takes the portfolio sensitivities as input (figure 2).
- (3) The expected funding cost of posting IM is called Margin Value Adjustment (MVA). Quantifying MVA requires the simulation of future IM (i.e. portfolio

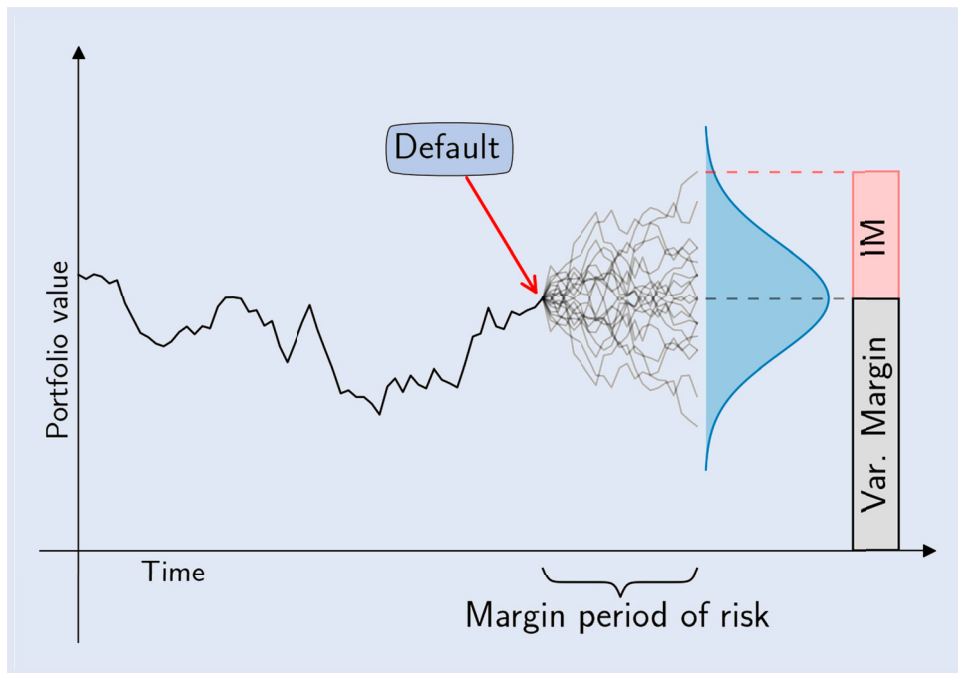


Figure 1. Definition initial margin (IM).




Figure 2. Standard initial margin model (SIMM).

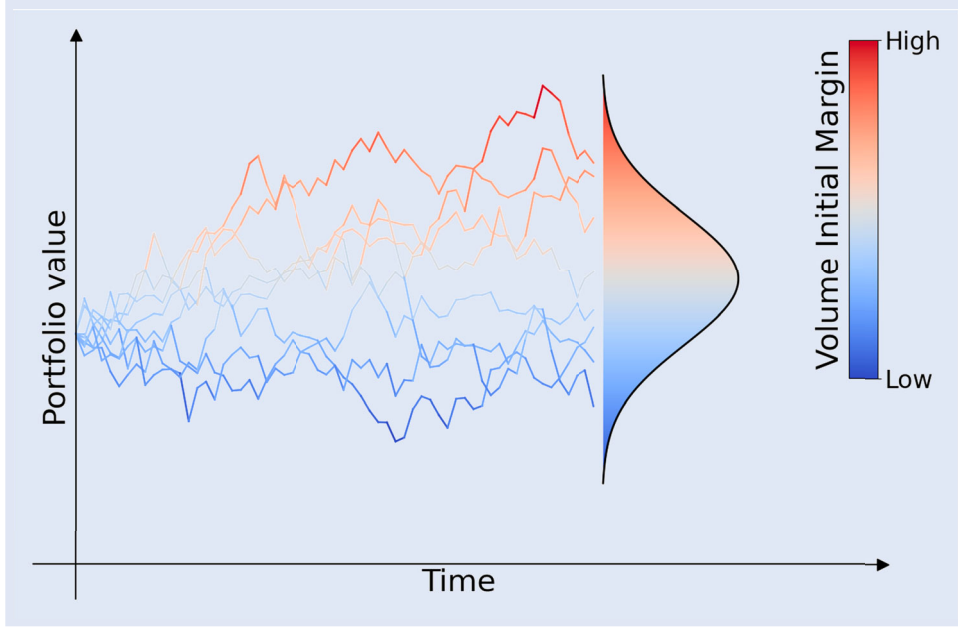


Figure 3. Definition margin value adjustment (MVA).

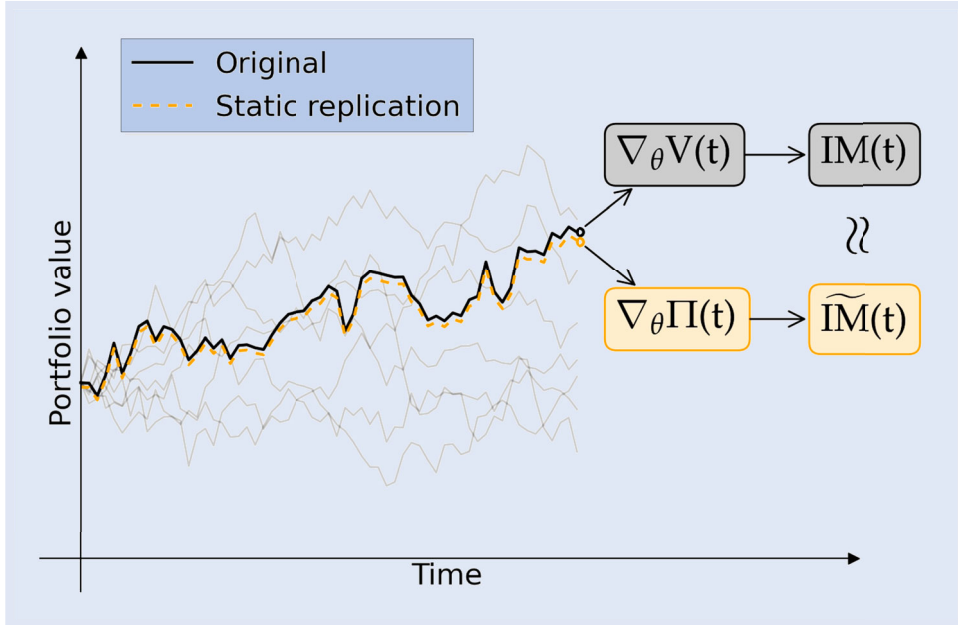


Figure 4. MVA through static replication.

sensitivities), which is a computational challenge (figure 3).

- (4) We propose a method to statically replicate a callable derivative with a portfolio of vanilla options. This simplifies the computation of path-wise sensitivities and future IM (figure 4).

2.2. Model and assumptions

First, we fix a finite time-horizon $[0, \bar{T}]$, on which we consider a continuous-time financial market. We assume the market is frictionless and free of arbitrage. We let the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ represent all possible states of the economy and the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, \bar{T}]}$ all information generated by the economy up to time- t .

We consider the notion of a *bank account* or *money market account*. Investments in the money market are assumed to compound a continuous, risk-free interest r_t . We refer to r_t as the *short-rate*. The time- t value of a unit of currency invested in the money market at time-zero is denoted as M_t , and we assume it satisfies the following dynamics

$$dM_t = r_t M_t dt, \quad M_0 = 1$$

We denote by \mathbb{Q} the risk-neutral measure equivalent to \mathbb{P} , which is associated with M_t as the numéraire. Attainable claims denominated by the numéraire are assumed to be martingales under \mathbb{Q} , which guarantees the absence of arbitrage (Harrison and Pliska 1981).

Throughout this paper we shall consider short-rate dynamics that are captured by an *affine term-structure model*, as

introduced in Duffie and Kan (1996). That means that the short-rate can be written as

$$r_t = g(x_t)$$

for some function $g : \mathbb{R}^d \rightarrow \mathbb{R}$. The vector $x_t := \{x_1(t), \dots, x_d(t)\}$ denotes a Markovian system of state-variables in \mathbb{R}^d . The state-variables are assumed to satisfy an SDE of the form

$$dx_t = \mu(t, x_t) dt + \eta(t) dW_t \quad (1)$$

with $W_t := \{W_1(t), \dots, W_d(t)\}$ denoting a d -dimensional Brownian motion under \mathbb{Q} and functions $\mu : [0, \bar{T}] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\eta : [0, \bar{T}] \rightarrow \mathbb{R}^{d \times d}$. We let $g(\cdot)$ and $\mu(t, \cdot)$ be affine functions of x_t , satisfying the standard regularity conditions, such that x_t admits a strong solution. Note that we restrict our scope to the Gaussian subclass of affine short-rate models. This is done to impose an intuitive relationship between model parameters and model-implied option volatilities (see section 3.5). This will facilitate the tractability of the computations; the estimation of Vega in particular. Generalizations to state-dependent diffusion coefficients should be possible but will involve the differentiation of Fourier transforms. This increases the complexity of the problem and falls outside the scope of this paper.

A zero-coupon bond is a contract that guarantees the holder one unit of currency at a pre-specified maturity date T . We will denote its time- t value as $P(t, T)$ for $t \leq T$. An important result for affine term-structure models, is that zero-coupon bond prices are exponential affine in x_t . See for example Andersen and Piterbarg (2010) and Duffie and Kan (1996) for details. Next to zero-coupon bonds, we will often refer to a related quantity, known as the *zero-rate*. We define a zero-rate as

$$R(t, T) := -\frac{\log P(t, T)}{\tau}$$

where τ denotes the year fraction between date t and T . For simplicity, we will assume that the collateral rate used for discounting and the instantaneous rate used to derive term-rates are both implied by the same short-rate r_t . In other words, our set-up will be a classic single-curve model environment. As term rates, we consider the classic, forward-looking LIBOR. Similarly one could consider a backward-looking, RFR-based term rate. In certain markets, the IBOR benchmark is discontinued, in favor of such *set-in-arrears* term rates (see Lyashenko and Mercurio 2019). As forward rates of forward- and backward-looking term rates are equivalent before the start of the accrual period, their treatment within this framework would mostly be the same.

2.3. Interest rate derivatives

The focus of this work will be on modeling Bermudan swaptions. Here we briefly introduce related derivatives and notation.

Interest rate swap: We consider fixed-for-floating interest rate swaps, where floating LIBOR payments are exchanged against fixed rate payments on dates T_0, \dots, T_M . The year-fraction between T_{j-1} and T_j will be denoted τ_j .

European swaption: A European swaption is a contract, which gives the holder the right, but not the obligation to enter an fixed-for-floating interest rate swap with pre-specified fixed rate K at a pre-specified future inception date T_0 . Considering the annuity $A_{0,M}(t) := \sum_{m=1}^M \tau_m P(t, T_m)$ as numéraire, the risk-neutral value of a swaption (see Filipovic 2009) can be expressed as

$$V_t = A_{0,M}(t) \mathbb{E}^{0,M} \left[\left(\delta (S_{0,M}(T_0) - K) \right)^+ \middle| \mathcal{F}_t \right] \quad (2)$$

where the expectation above is taken under the annuity measure $\mathbb{Q}^{0,M}$. The parameter δ distinguishes between a payer ($\delta = 1$) and a receiver ($\delta = -1$) swaption. Under the assumption that the swap rate has Gaussian dynamics, the swaption price can be evaluated using Bachelier's formula (see Musiela and Rutkowski 1997). Given a swaption has market price $V_{\text{mkt}}(t, K)$, the *implied normal volatility* is defined as the unique scalar σ_{imp} that solves the equation

$$V_{\text{Bach}}(t, \sigma_{\text{imp}}, K) = V_{\text{mkt}}(t, K)$$

Bermudan swaption: A Bermudan swaption is a contract, which gives the holder the right to enter a fixed-for-floating interest rate swap with maturity T_M at any of the dates T_0, \dots, T_{M-1} at a pre-specified fixed rate K . Should the holder of the Bermudan choose to exercise at T_m , the payoff is given by

$$\begin{aligned} h(T_m, x_{T_m}) &:= h_m(x_{T_m}) \\ &= \delta A_{m,M}(T_m; x_{T_m}) [S_{m,M}(T_m; x_{T_m}) - K] \end{aligned}$$

The value of a Bermudan swaption can be represented through a dynamic programming formulation, in which the price is captured by means of a backward recursion (Glasserman 2004), given by

$$V(t, x_t) = \begin{cases} \max\{h_{M-1}(x_t), 0\} & \text{if } t = T_{M-1} \\ \max\{h_m(x_t), C_m(t, x_t)\} & \text{if } t = T_m, \\ & m = 0, \dots, M-2 \\ C_{m-1}(t, x_t) & \text{if } t \in (T_{m-1}, T_m) \end{cases}$$

where C_m represents the *continuation value* of the contract.

2.4. Credit risk metrics

In this work, we will consider credit exposure and forward sensitivity profiles of Bermudan swaptions. Exposure is essential for CVA computation. Product sensitivities, Delta and Vega in particular, serve as an input to ISDA SIMM calculations for initial margin.

2.4.1. Exposure and CVA. The *expected positive exposure* (EPE) of a financial contract with price V_t at time $t \geq 0$ is

$$\text{EPE}(t) := \mathbb{E}^{\mathbb{Q}} \left[e^{-\int_0^t r_u du} \max\{V_t, 0\} \middle| \mathcal{F}_0 \right]$$

Expected positive exposure represents the expected loss on a claim in the event of a counterparty defaulting at a given time

$t \geq 0$ (Gregory 2020). The expectation is evaluated under the risk-neutral measure \mathbb{Q} . Exposure is a key ingredient in the quantification of counterparty risk at trade or counterparty level through CVA. CVA is the difference between the total value of a derivative in a market that is completed with counterparty risk and its default-free value (Burgard and Kjaer 2013). Let $\tau \geq 0$ denote the default time of the counterparty. Assuming τ is independent of exposure (i.e. ignoring wrong- or right-way risk), CVA can be computed as follows Green (2015)

$$\begin{aligned} \text{CVA} &= (1 - \text{RR}) \int_0^T \text{EPE}(u) d\mathbb{Q}(\tau = u) \\ &\approx \text{LGD} \sum_{i=1}^n \text{EPE}(t_i) \times \text{PD}(t_{i-1}, t_i) \end{aligned} \quad (3)$$

with RR denoting the recovery rate, LGD the loss given default and PD the probability of default of the counterparty. A common approach in the literature is to model the default time τ as an \mathcal{F}_t -adapted stopping time using a Cox-construction (Jeanblanc and Li 2020). In that case, the survival function of the counterparty follows an exponential distribution, i.e. $\mathbb{Q}(\tau > t) = \exp\{-\Lambda(t)\}$ with $\Lambda(t)$ denoting a hazard function. The hazard function can be specified in terms of a hazard rate $\lambda(t)$, which yields $\Lambda(t) = \int_0^t \lambda(u) du$. Therefore, the probability of default can be written as

$$\text{PD}(t_{i-1}, t_i) = e^{-\int_0^{t_{i-1}} \lambda(u) du} - e^{-\int_0^{t_i} \lambda(u) du} \quad (4)$$

2.4.2. Initial margin and MVA. An important means to reduce the exposure at risk is the exchange of collateral. In this work, we will consider *initial margin* (IM) and its related funding costs. IM is a protection against exposure changes during the margin period of risk, the time interval after a default when outstanding positions are not yet settled, but VM is no longer updated (Green and Kenyon 2015). Typically, IM reflects the 99% value-at-risk of the portfolio change over a 10-day time interval. Throughout this work, we will consider the *Standard Initial Margin Model* (SIMM) to be the default method for IM computations. The main take-away from SIMM-IM is that it is a sensitivity-based approach. The idea is that the response of an instrument to a risk-factor shock is efficiently approximated by multiplying the instrument's risk-factor sensitivity by the corresponding shock size (International Swaps and Derivatives Association, Inc 2016). We will provide a brief summary of SIMM below, but refer to International Swaps and Derivatives Association, Inc (2020) for details.

Consider risk-factors $\theta := (\theta_1, \dots, \theta_K)^\top$ and denote the 10-day risk-factor increments as $\Delta\theta_1, \dots, \Delta\theta_K$. The variance of the total response of an instrument V_t to each shock $\Delta\theta_i$ is computed as

$$\text{Var} \left(\sum_{i=1}^K \frac{\partial V_t}{\partial \theta_i} \Delta\theta_i \right) = [D_\theta V_t]^\top \Sigma_\theta [D_\theta V_t]$$

where $D_\theta V_t$ denotes the gradient of V_t w.r.t. θ and Σ_θ denotes the covariance matrix of θ . If the total response defined above,

is assumed to be approximately Gaussian with mean zero, the VaR can be estimated as follows

$$\text{VaR}_{0.99}(\Delta V_t) \approx \Phi^{-1}(0.99) \sqrt{[D_\theta V_t]^\top \Sigma_\theta [D_\theta V_t]}$$

Where Φ^{-1} denotes the inverse of the standard normal CDF. For interest rate derivatives, ISDA distinguishes three relevant types of risk-factors, namely *Delta-risk*, *Vega-risk* and *Curvature-risk*, the latter referring to the second-order Gamma impact. The total IM is given by

$$\text{IM} := \text{Margin}_{\text{Delta}} + \text{Margin}_{\text{Vega}} + \text{Margin}_{\text{Curvature}} \quad (5)$$

where each component represents a VaR estimation w.r.t. a type of underlying risk-factors, i.e. market interest rates and market implied volatilities. The risk-factors of each type are subdivided into twelve buckets, corresponding to the yields at the tenors $\mathcal{T} := \{2W, 1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 10Y, 15Y, 20Y, 30Y\}$. Hence for each component one needs to compute the corresponding 12-dimensional sensitivity vector, which are specified as (see also International Swaps and Derivatives Association, Inc 2020):

Delta : The PV01 of the instrument V_t w.r.t. the risk-free interest rate R_k at tenor $\tau_k \in \mathcal{T}$, i.e. $\Delta_k = V_t(R_k + 1\text{bp}) - V_t(R_k)$.

Vega : The sensitivity of the instrument V_t w.r.t. volatility risk-factor σ_k , i.e. $\mathcal{V}_k = V_t(\sigma_k + 1\%) - V_t(\sigma_k)$. Here, σ_k denotes the implied at-the-money volatility of a swaption with an expiry time equal to $\tau_k \in \mathcal{T}$.

Curvature : Curvature margin represents the second-order impact of the interest rate risk-factors to the VaR estimate. In SIMM, the Gamma is approximated with a Vega-Gamma relationship. Therefore, the computation of Curvature margin takes the Vega sensitivity vector as input, similar to Vega margin.

For SIMM-IM quantification in practice, the user only needs to compute the portfolio sensitivities. The other parameters are provided by ISDA in an annual publication.

A key difference between VM and IM is that the former is symmetric. VM reflects the value of the underlying trade and can be rehypothecated. This means the receiver can reuse the collateral, for example in the margin agreement of an opposite transaction. In contrast, IM reflects the risk of a trade and must be exchanged in both directions, without netting the amounts. The collateral is often required to be posted on a segregated account and therefore cannot be rehypothecated (Gregory 2020). As a result, IM must be funded by the dealer, which comes at a cost. The expected lifetime cost of posting IM against a portfolio is called *Margin Value Adjustment* (MVA) (Green and Kenyon 2015).

The size of MVA is driven by the dynamic IM profile of the portfolio and the funding spread that applies to the IM posting. As the volume of IM is dependent on the state of the economy, $\{\text{IM}(t) : t \in [0, T]\}$ represents a stochastic process. The *expected initial margin* (EIM) of a financial contract at

time $t \geq 0$ is defined as

$$\text{EIM}(t) := \mathbb{E}^{\mathbb{Q}} \left[e^{-\int_0^t r_u du} \text{IM}(t) \middle| \mathcal{F}_0 \right]$$

EIM represents the expected volume of IM that needs to be posted by the dealer at a given time $t \geq 0$ (Gregory 2020). With this expression at hand, MVA can be computed as follows

$$\begin{aligned} \text{MVA} &= \int_0^T \text{FS}(u) \text{EIM}(u) du \\ &\approx \sum_{i=1}^n \text{FS}(t_i) \times \text{EIM}(t_i) \times (t_i - t_{i-1}) \end{aligned} \quad (6)$$

with FS denoting the funding spread, which reflects the spread on the collateral rate w.r.t. the risk-free rate (Green and Kenyon 2015).

3. The static replication method

In this section, we describe the algorithm that is used to compose a static replication portfolio for a Bermudan swaption. The algorithm is embedded in a Monte Carlo framework and relies on a regress-later technique. See Glasserman and Yu (2004) for a general introduction to regress-later. The method is strongly inspired by the work of Jain and Oosterlee (2015) and extends the approach that is presented in Lokeshwar *et al.* (2022) and Hoencamp *et al.* (2023).

In Jain and Oosterlee (2015), it is shown that early-exercise options can be priced with an application of a regress-later technique. As such, the ‘later’ option price $V_{T_{m+1}}$ is regressed against the ‘later’ risk-factor realization $x_{T_{m+1}}$. The continuation value is subsequently estimated by evaluating the conditional expectation of the regressed option value, i.e.

$$\begin{aligned} C_m(T_m) &= M_{T_m} \mathbb{E}^{\mathbb{Q}} \left[\frac{V_{T_{m+1}}}{M_{T_{m+1}}} \middle| \mathcal{F}_{T_m} \right] \\ &\approx M_{T_m} \mathbb{E}^{\mathbb{Q}} \left[\frac{\sum_{k=1}^K \alpha_k \phi_k(x_{T_{m+1}})}{M_{T_{m+1}}} \middle| \mathcal{F}_{T_m} \right] \end{aligned}$$

which can be computed in (semi) closed-form for an appropriate choice of basis-functions ϕ_k .

In Lokeshwar *et al.* (2022), it is proposed to perform the regression with a shallow neural network (i.e. a feed-forward neural network with a single hidden layer). By considering the appropriate network structure, it is shown that the regression can be interpreted as a portfolio of short-maturity options. As a consequence, the continuation value can be evaluated by simply pricing the replication instruments. In Hoencamp *et al.* (2023), this approach is extended to an interest rate modeling framework. There it is shown that a Bermudan swaption can be semi-statically replicated by an options portfolio written on zero-coupon bonds. The central concept of the latter two studies is that if a portfolio perfectly reproduces the value function of a derivative security at some future time T_j for every realization of the state-variables, the no-arbitrage

condition implies that this portfolio will also replicate the security at any time $t < T_j$, as long as no cash-flows can occur between t and T_j . Where a dynamic replicating portfolio needs to be rebalanced continuously through time, the semi-static replication only needs rebalancing on a finite number of instances.

The algorithm that is proposed in this paper builds further on the foundations that were laid by Jain and Oosterlee (2015), Lokeshwar *et al.* (2022) and Hoencamp *et al.* (2023). We suggest two main novelties compared to the earlier studies

- We propose to use the swap rate as the regression variable, which is an implicit risk-factor in our model set-up. As a consequence, the regression can be interpreted as the payoff of a portfolio of European swaptions. We further elaborate on this in Section 3.4.
- We propose a new variation to the algorithm, which allows us to compose a replication that is not semi-static, but in fact truly static. In other words, this algorithm lets one compose a portfolio at time zero, which will mirror the value of the Bermudan until it is either exercised or expired, without the need of updating the portfolio at intermediate time-points. We further elaborate on this in Section 3.3.

3.1. The algorithm

The algorithm is executed in an iterative manner, moving backwards in time. Starting at the final exercise date T_{M-1} , a regression is performed to determine the weights and strikes of an options portfolio which replicates the Bermudan’s payoff. This replication is subsequently used to estimate the value of the Bermudan at the preceding monitor date, yielding the target function of the consecutive regression. This process is repeated until the first monitor date T_0 and a portfolio has been composed which statically replicates the Bermudan between time-zero and its final exercise date. Below we will describe the algorithm in more detail.

Start by sampling N trajectories of the underlying state-variable. For each realization of $x_{T_{M-1}}$, compute the Bermudan swaption payoff at T_{M-1} (the final exercise opportunity) $f_{M-1}(x_{T_{M-1}}) = \max\{h_{M-1}(x_{T_{M-1}}), 0\}$. This will act as the target function of the first regression. Now recursively execute the following steps for $m = M - 1, \dots, 0$:

1. Select and sample the regression variable

Select an asset $z_m(t) = \{z_{m,1}(t, x_t), \dots, z_{m,d}(t, x_t)\}$ taking value in \mathbb{R}^d that will act as the independent regression variable. For compactness, we will suppress the dependency of z_m on the state variable x_t , when clear from the context. The asset(s) should satisfy the following conditions:

- $z_m(T_m)$ is a square integrable random variable on the probability space $(\Omega, \mathcal{F}_{T_m}, \mathbb{Q})$.
- For each $m = 0, \dots, M - 1$, the realization $f_m(x_{T_m})$ is uniquely determined by $z_m(T_m)$.
- The parametrization $\{(z_m(T_m; x), f_m(x)) : x \in \mathbb{R}^d\}$ is continuous.

where f_m denotes a target function, which we define in (8) below. Subsequently compute for each sample path the realizations of $z_m(T_m)$.

2. Regress the target function

Consider a regression function $G_m : \mathbb{R}^d \rightarrow \mathbb{R}$ that is defined as

$$G_m = \psi_2 \circ A_2 \circ \psi_1 \circ A_1 \quad (7)$$

Here we denote by $A_1 : \mathbb{R}^d \rightarrow \mathbb{R}^n$, $A_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ affine functions of the form

$$A_1(y) \mapsto \mathbf{w}_1 y + \mathbf{b}_1 \quad \text{and} \quad A_2(y) \mapsto \mathbf{w}_2 y + b_2$$

with $\mathbf{w}_1 \in \mathbb{R}^{d \times n}$, $\mathbf{w}_2 \in \mathbb{R}^{n \times 1}$, $\mathbf{b}_1 \in \mathbb{R}^n$ and $b_2 \in \mathbb{R}$. Furthermore we denote by $\psi_1 : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $\psi_2 : \mathbb{R} \rightarrow \mathbb{R}$ activation functions of the form

$$\psi_1(y) \mapsto (\max\{y_1, 0\}, \dots, \max\{y_n, 0\}) \quad \text{and} \quad \psi_2(y) \mapsto y$$

Additionally, consider a target function $f_m : \mathbb{R}^d \rightarrow \mathbb{R}$ of the form

$$f_m(x) = \begin{cases} q_m(x) \cdot \max\{h_m(x), 0\} & \text{if } m = M - 1 \\ q_m(x) \cdot \max\{h_m(x) - \tilde{C}_m(T_m, x), 0\} & \text{if } m = M - 2, \dots, 0 \end{cases} \quad (8)$$

where $q_m(x)$ denotes a possible scaling factor. The rationale behind this target function is that we aim to regress the fraction of the option value that is lost, when an exercise date passes, but the option is not exercised. We provide an elaborate motivation and interpretation in section 3.3.

Subsequently, optimize the parameters $\beta_m = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{b}_1, b_2\}$, such that $G_m(z_m(T_m); \beta_m)$ fits the target function $f_m(x_{T_m})$. The parameters are determined by minimizing the cost function

$$L_m(\beta) = \mathbb{E}^{\mathbb{Q}} \left[\left(G_m(z_m(T_m); \beta) - f_m(x_{T_m}) \right)^2 \right] \quad (9)$$

$$\approx \frac{1}{N} \sum_{n=1}^N \left(G_m(z_m(n); \beta) - f_m(x_{T_m}(n)) \right)^2 \quad (10)$$

where $z_m(n)$ and $x_{T_m}(n)$ denote sampled realizations of the asset and the state variable, respectively, on a corresponding MC path.

3. Estimate the continuation value

Finally, the continuation value of the Bermudan swaption at the preceding exercise date T_{m-1} is estimated by calculating

$$\tilde{C}_{m-1}(T_{m-1}, x_{T_{m-1}}) = \sum_{j=m}^{M-1} M_{T_{m-1}} \mathbb{E}^{\mathbb{Q}} \left[\frac{G_j(z_j(T_j); \beta_j)}{M_{T_j} \cdot q_j(x_{T_j})} \middle| x_{T_{m-1}} \right] \quad (11)$$

With an appropriate choice of z_m , this can be computed in closed-form or approximated with an efficient numerical routine.

3.2. Interpretation of the regression function

The regression function as presented in (7) can be considered a shallow neural network, i.e. a feed-forward neural network with a single hidden layer. The structure is graphically represented in figure 5(a). The outcome of the first hidden layer is a vector in \mathbb{R}^n , consisting of n neurons. Each individual neuron $O_{m,i}$ is calculated as

$$O_{m,i} = \max \left\{ \sum_{j=1}^d w_{ij} z_{m,i}(T_m) + b_i, 0 \right\}$$

The functional form of each hidden node thus corresponds to the payoff of an arithmetic basket option written on $z_{m,1}, \dots, z_{m,d}$.

The outcome of the second layer is calculated as

$$\sum_{i=1}^n w_i O_{m,i} + b$$

The functional form of the outcome of the neural network, hence corresponds to the payoff of a weighted portfolio of basket options, plus a zero-coupon bond maturing at T_m . Computing the conditional expectation given in (11) can be interpreted as evaluating the risk-neutral price of this portfolio.

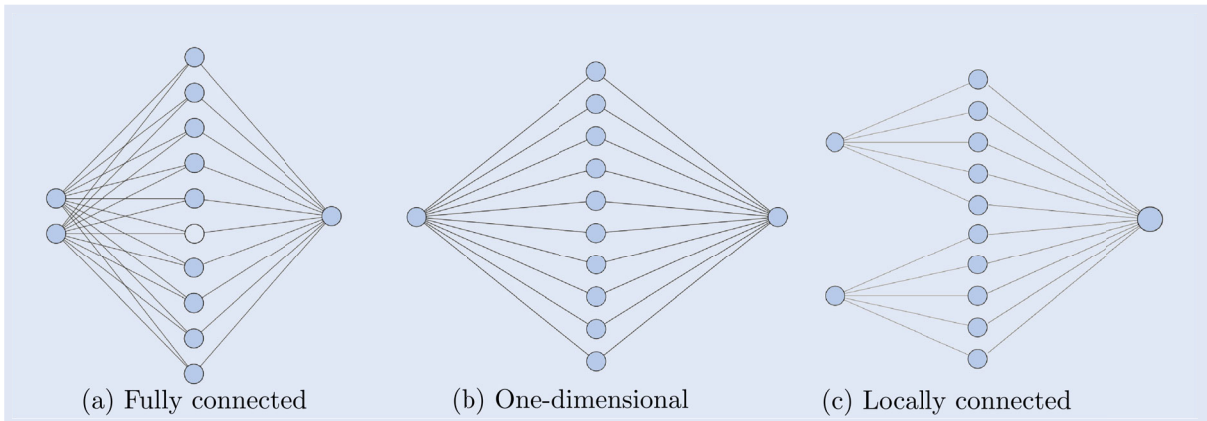


Figure 5. Neural network structures. (a) Fully connected. (b) One-dimensional and (c) Locally connected.

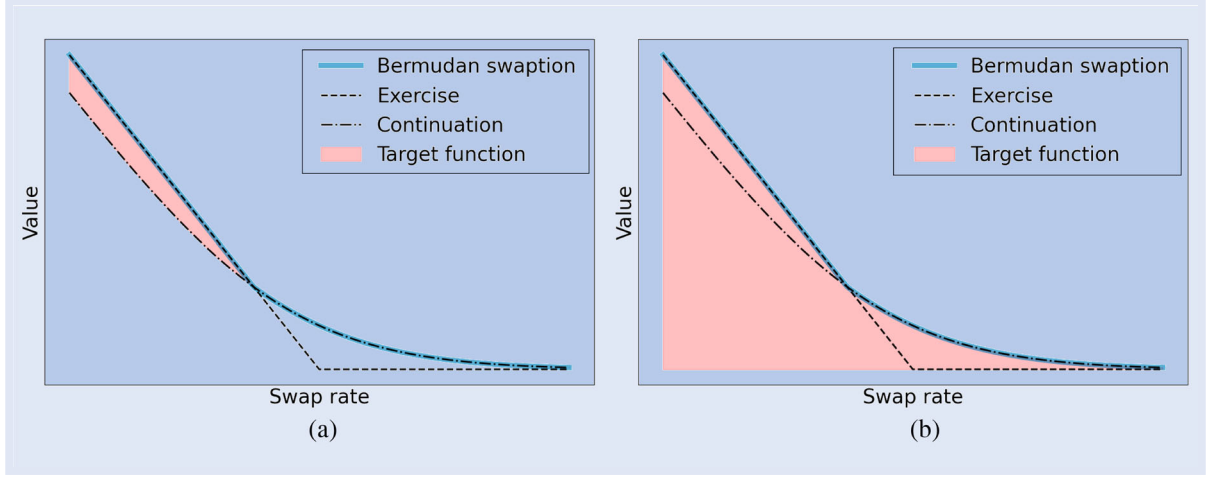


Figure 6. Target functions regress-later: two approaches. (a) Fully static replication and (b) Semi-static replication.

3.3. A static replication

Let $\Pi_m(t)$ denote an options portfolio maturing at T_m with payoff equal to

$$\Pi_m(T_m) = \frac{1}{q_m(x_{T_m})} G_m(z_m(T_m); \beta_m) \quad (12)$$

By the deliberate choice of the target function f_m in (8), the payoff of this portfolio yields the approximation

$$\begin{aligned} \Pi_m(T_m) &\approx \max \left\{ h_m(x_{T_m}) - \tilde{C}_m(T_m, x_{T_m}), 0 \right\} \\ &= \max \left\{ h_m(x_{T_m}), \tilde{C}_m(T_m, x_{T_m}) \right\} - \tilde{C}_m(T_m, x_{T_m}) \end{aligned}$$

which is graphically represented by the red area in figure 6(a). As a consequence of the no-arbitrage condition, it follows that at any time $t < T_m$ the risk-neutral portfolio value yields the approximation $\Pi_m(t) \approx C_{m-1}(t, x_t) - C_m(t, x_t)$. By a repeated argument, whenever $t \in (T_{m-1}, T_m)$ we have

$$\sum_{i=m}^{M-1} \Pi_i(t) \approx C_{m-1}(t, x_t) = V(t)$$

The combined portfolio $\Pi(t) := \sum_{i=0}^{M-1} \Pi_i(t)$, composed at $t = 0$, can be interpreted as a fully static replication of the Bermudan swaption. Whenever a monitor date T_m is reached, subportfolio Π_m expires. Then, two things could happen:

- $C(T_m, x_{T_m}) \leq h_m(x_{T_m})$: The holder exercises the Bermudan and the value of $\sum_{i=m+1}^{M-1} \Pi_i$ plus the payoff of Π_m will equal the exercise value.
- $C(T_m, x_{T_m}) > h_m(x_{T_m})$: The holder holds the Bermudan, subportfolio Π_m expires worthless and $\sum_{i=m+1}^{M-1} \Pi_i$ will continue to replicate the Bermudan.

Thus, portfolio Π will mirror the Bermudan swaption value until it is either matured or exercised.

The algorithm presented in this work is a variation to the semi-static replication approach presented in Lokeshwar *et al.* (2022) and Hoencamp *et al.* (2023). There, the target function f_m is selected such that the portfolio payoff Π_m yields the approximation

$$\Pi_m(T_m) \approx \max \left\{ h_m(x_{T_m}), \tilde{C}_m(T_m, x_{T_m}) \right\}$$

which is graphically represented by the red area in figure 6(b). In that case, subportfolio Π_m is set-up on T_{m-1} , it replicates the Bermudan on the time-interval (T_{m-1}, T_m) and expires at T_m . Its payoff will equal the exercise value of the Bermudan, or suffice to set-up subportfolio Π_{m+1} in case the Bermudan is continued. The advantage of the latter approach is that only $\Pi_m(t)$ needs to be priced in order to value the Bermudan, rather than $\sum_{i=m}^{M-1} \Pi_i(t)$. A disadvantage is that the replication is only semi-static and needs rebalancing on each monitor date T_m .

3.4. Estimation of the continuation value

How the continuation value as given in (11) is computed, depends on the specification of the model and the selection of the regression asset. In the general case, the portfolio Π_m can be interpreted as a portfolio of arithmetic basket options. The valuation of a basket option is typically still not an easy exercise. However, with the appropriate constraints on the neural network structure used for the regression, the replicating portfolio can be reduced to a set of European options. A first approach would be to consider a one-dimensional regression asset, which corresponds to the design depicted in figure 5(b). This would be appropriate under a one-factor model. A second approach would be to constrain matrix A_1 to have only one non-zero entry in each row. This corresponds to the neural network design depicted in figure 5(c). In the numerical examples included in this work, we will be considering a 1-factor model environment. Generalizations to multi-factor models have been described in Hoencamp *et al.* (2023). We will consider the following assumptions:

- *Regression asset*: select $z_m(t) = S_{m,M}(t)$.
- *Scaling factor*: set $q_m(x_t) = \frac{1}{A_{m,M}(t, x_t)}$.

Under the assumptions stated above, the (sub)portfolio Π_m as defined in (12), can be interpreted as a portfolio of European swaptions, written on the swap rate $S_{m,M}(t)$.[†] The

[†] For simplicity we will ignore the zero-coupon bond in the replicating portfolio (see section 3.2). This is achieved by constraining $b_2 = 0$. Our numerical experiments indicate a minimal impact on the accuracy of the replication.

valuation of $\Pi_m(t)$ for some $t < T_m$ comes down to

$$\begin{aligned} \Pi_m(t) &= M_t \mathbb{E}^{\mathbb{Q}} \left[\frac{G_m(z_m(T_m); \beta_j)}{M_{T_m} \cdot q_m(x_{T_m})} \middle| x_t \right] \\ &= M_t \mathbb{E}^{\mathbb{Q}} \left[\frac{\sum_{i=1}^n w_{2,i} \max\{w_{1,i} S_{m,M}(T_m) + b_i, 0\}}{M_{T_j}} \middle| x_t \right] \\ &= \sum_{i=1}^n w_{2,i} A_{m,M}(t) \mathbb{E}^{m,M} \left[\max\{w_{1,i} S_{m,M}(T_m) + b_i, 0\} \middle| x_t \right] \end{aligned}$$

where the expectation in the last line is taken under the annuity measure $\mathbb{Q}^{m,M}$ associated to the numéraire $A_{m,M}(t)$. The expression above should be recognized as a linear combination of European swaption prices (see equation (2)). The parameter $w_{1,i}$ distinguishes between a payer ($w_{1,i} > 0$) and receiver ($w_{1,i} < 0$) swaption. The parameter $w_{2,i}$ distinguishes between a long ($w_{2,i} > 0$) and a short ($w_{2,i} < 0$) position in the swaption contract.

3.5. Valuation of the replicating portfolio

Under the assumptions of section 3.4, portfolio Π_m consists of European swaptions, which means the valuation comes down to pricing each individual option. We apply a coefficient freezing technique, described in Andersen and Piterbarg (2010) and Schrager and Pelsser (2006). With this technique, the swap rate process is approximated as a generalized arithmetic Brownian motion, by freezing the stochastic terms in its diffusion. The option's implied volatility can then be approximated by integrating the 'frozen' diffusion coefficient. We briefly summarize the method below, and refer to Andersen and Piterbarg (2010) and Schrager and Pelsser (2006) for details.

Consider the dynamics of the swap rate. By Itô's lemma, it can be established that $S_{m,M}$ satisfies the SDE

$$dS_{m,M}(t) = (\dots) dt + [D_x S_{m,M}(t)]^\top \eta(t) \cdot dW_t$$

Considering the annuity as numéraire, the swap rate can be regarded as the quotient of some tradable assets and the numéraire itself. As a consequence, the swap rate is a Martingale under the annuity measure $\mathbb{Q}^{m,M}$ (Jamshidian 1997). Hence, for $0 \leq t \leq T \leq T_m$ we can write

$$S_{m,M}(T) = S_{m,M}(t) + \int_t^T \Sigma_{m,M}(u)^\top \eta(u) \cdot dW_u^{m,M} \quad (13)$$

where $W^{m,M}$ denotes a d -dimensional $\mathbb{Q}^{m,M}$ Brownian motion and $\Sigma_{m,M}(t) := [D_x S_{m,M}(t)]^\top$. Working out the diffusion for $i = 1, \dots, d$ yields

$$\begin{aligned} \Sigma_{m,M}(t) &= B(t, T_m) \frac{P(t, T_m)}{A_{m,M}(t)} - B(t, T_m) \frac{P(t, T_m)}{A_{m,M}(t)} \\ &\quad - S_{m,M}(t) \sum_{l=m+1}^M \tau_l B(t, T_l) \frac{P(t, T_l)}{A_{m,M}(t)} \end{aligned} \quad (14)$$

The diffusion coefficient is stochastic, due to the $\frac{P(t, T_l)}{A_{m,M}(t)}$ terms. However, these terms again represent the quotient of a tradable asset and the numéraire. It is conjectured in Schrager and Pelsser (2006), that these terms are Martingales with low quadratic variation. A good approximation is therefore achieved by freezing them at their time- t value, resulting in a deterministic diffusion coefficient $\tilde{\Sigma}_{m,M}(u)$ for $u \in [t, T]$. It follows

$$S_{m,M}(T) \approx \tilde{S}_{m,M}(T) := S_{m,M}(t) + \int_t^T \tilde{\Sigma}_{m,M}(u)^\top \eta(u) \cdot dW_u^{m,M} \quad (15)$$

The approximation $\tilde{S}_{m,M}(T)$ conditioned on \mathcal{F}_t is normally distributed. Applying Itô isometry allows us to approximate the implied volatility of an option written on $S_{m,M}$, which yields

$$\begin{aligned} \tilde{\sigma}_{\text{imp}}(t, T) &= \sqrt{\frac{1}{T-t} \mathbb{E}^{m,M} \left[\left\| \int_t^T \tilde{\Sigma}_{m,M}(u)^\top \eta(u) \cdot dW_u^{m,M} \right\|_2^2 \middle| x_t \right]} \\ &= \sqrt{\frac{1}{T-t} \int_t^T \|\tilde{\Sigma}_{m,M}(u)^\top \eta(u)\|_2^2 du} \end{aligned} \quad (16)$$

$$= \sqrt{\frac{1}{T-t} \int_t^T \|\tilde{\Sigma}_{m,M}(u)^\top \eta(u)\|_2^2 du} \quad (17)$$

Substituting $\tilde{\sigma}_{\text{imp}}(t, T)$ into Bachelier's formula results in the desired swaption price.

4. Numerical algorithms and mathematical consistency

In this section, we describe the numerical routines for computing sensitivities along the Monte Carlo path. In short, we show how to compute the sensitivities of the replicating portfolio in a time-consistent manner. This involves the differentiation of the replicating instrument values (i.e. greeking) and the differentiation of the portfolio weights.

Our approach to the latter is inspired by the work of Jain et al. (2019). In Jain et al. (2019) a regress-later approach is considered, which utilizes ordinary least-square regression. As such, the weights of the basis-functions can be obtained explicitly as a function of the risk-factor and consistently differentiated. In this paper, only an implicit relation between the risk-factors and the portfolio weights is known through the minimized cost function (equation (10)). Yet, we show that an estimator of the weight sensitivities can still be obtained through an application of the implicit function theorem.

An advantage of our approach is that it should be easier to evaluate sensitivities to implicit risk-factors, which are not directly modeled. In the context of MVA this concerns the sensitivities to the realizations of future market rates and future implied volatilities. With regard to static replication, the replicating portfolio is selected such that the required future Delta and Vega are tractable to compute. In contrast, with, for example, an application of AD

on LSM, this may not be straight-forward. It would require the computation pseudo Jacobian-inverses or adding plenty pseudo-interpolation-nodes to the term-structure of zero-rates and instantaneous volatilities, which may be a tedious task.

4.1. Notation and assumptions

Let $t \in [0, \bar{T})$ and let V_t denote the time- t risk-neutral value of a Bermudan swaption with strike K and tenor structure T_0, T_1, \dots, T_M , conditioned on the event that the option is not yet exercised at t . Next to the Bermudan swaption, we consider a static replicating portfolio Π consisting of European swaptions. We write

$$\Pi(t) = \begin{cases} \sum_{i=0}^{M-1} \Pi_i(t) & \text{if } 0 \leq t \leq T_0 \\ \sum_{i=m}^{M-1} \Pi_i(t) & \text{if } T_{m-1} < t \leq T_m \\ 0 & \text{if } T_{M-1} < t \end{cases}$$

where Π_m denotes a subportfolio of swaptions, written on the swap rate $S_{m,M}$. The subportfolios are defined as

$$\begin{aligned} \Pi_m(t) &:= \mathbf{w}_m^\top \boldsymbol{\pi}_m(t) = \sum_{i=1}^n w_{m,i} \pi_{m,i}(t) \\ \pi_{m,i}(t) &= A_{m,M}(t) \mathbb{E}^{m,M} \left[\left(\delta_{m,i} (S_{m,M}(T_m) - K_{m,i}) \right)^+ \middle| x_t \right] \end{aligned}$$

Where we denote $\boldsymbol{\pi}_m(t) := (\pi_{m,1}(t), \dots, \pi_{m,n}(t))^\top$, $\mathbf{w}_m := (w_{m,1}, \dots, w_{m,n})^\top$ and $\mathbf{k}_m := (K_{m,1}, \dots, K_{m,n})^\top$

To keep computations feasible, we will consider two assumptions:

- The time-zero zero-rate curve $\tau \mapsto R(0, \tau)$ is assumed to be an interpolation between nodes $\{R_0, R_1, \dots, R_K\}$, with $R_k := R(0, \tau_k)$. The nodes are assumed to be the result of a calibration bootstrap against market rates, such as swap- or deposit rates.
- The instantaneous volatility $\eta(t)$ is assumed to be piece-wise constant between dates $\{\tau_0, \tau_1, \dots, \tau_K\}$. The volatility level η_k between τ_{k-1} and τ_k , is assumed to be the result of a bootstrap against market implied volatilities σ_k of European swaptions with maturity τ_k .

4.2. Sensitivities along the monte carlo path

In this work, we focus on two types of sensitivities, namely *Delta* and *Vega*. The time-zero sensitivities are specified as follows: Let $\Delta := (\Delta_0, \Delta_1, \dots, \Delta_K)$, where Δ_k denotes the sensitivity of V_0 to an infinitesimal bump of the zero-rate node R_k . It is defined as

$$\Delta_k := \lim_{\varepsilon \rightarrow 0} \frac{V_0(R_k + \varepsilon) - V_0(R_k)}{\varepsilon}$$

The impact of a bump in a zero-rate node, while keeping other nodes fixed, is graphically represented in figure 7(a). Let $\mathcal{V} := (\mathcal{V}_0, \mathcal{V}_1, \dots, \mathcal{V}_K)$, where \mathcal{V}_k denotes the sensitivity of V_0 to an infinitesimal bump of implied volatility σ_k corresponding to the swaption with maturity τ_k . It is defined as

$$\mathcal{V}_k := \lim_{\varepsilon \rightarrow 0} \frac{V_0(\sigma_k + \varepsilon) - V_0(\sigma_k)}{\varepsilon}$$

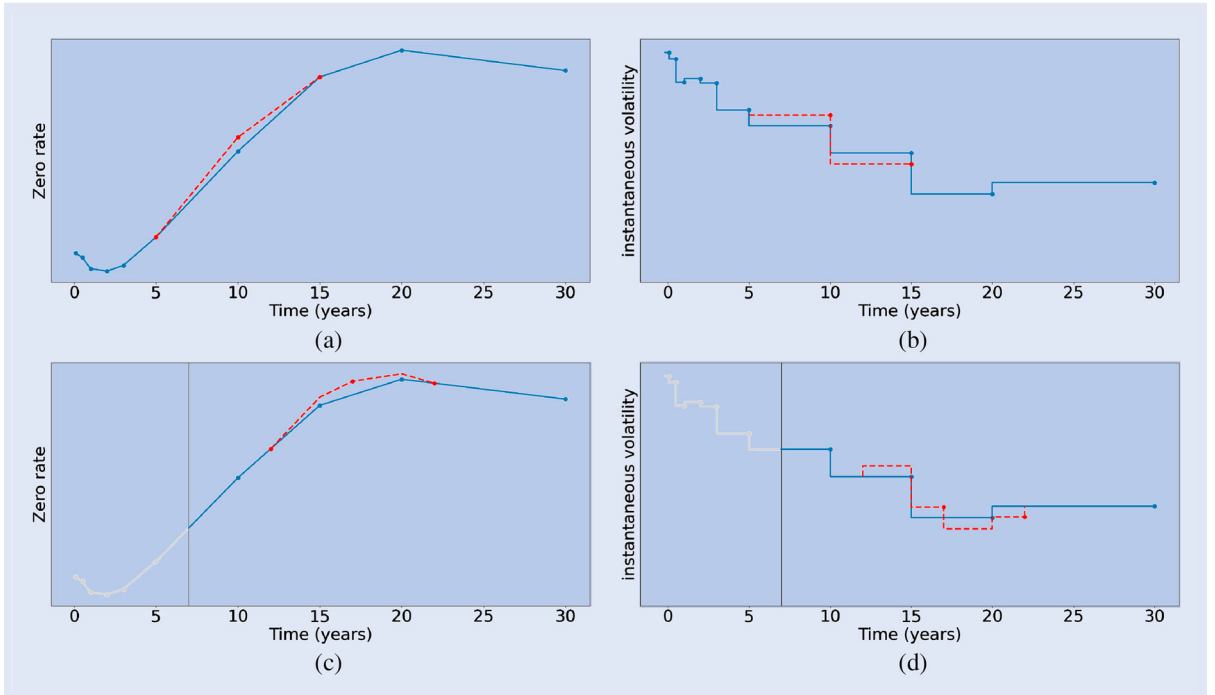


Figure 7. Impact of bumping the market rates (left) and implied volatilities (right) at t -zero (top) and a future simulation date (bottom). All bumps correspond to bucket $\tau_k = 10$. (a) Zero-rate bump at $t = 0$. (b) Instant. volatility bump $t = 0$. (c) Zero-rate bump at $t = 7$ and (d) Instant. volatility bump at $t = 7$.

A bump in σ_k , while keeping other implied volatilities fixed will impact the level of η_k and η_{k+1} . This impact is graphically represented in figure 7(b).

The sensitivities along the Monte Carlo path are understood as future-time generalizations of the quantities defined above. That is, future Delta $\Delta_k(t)$ is interpreted as the sensitivity of V_t to an infinitesimal bump of the hypothetical zero-rate node $R_k(t) := R(t, t + \tau_k)$, conditional on \mathcal{F}_t . This is graphically represented in figure 7(c). Future Vega $\mathcal{V}_k(t)$ is interpreted as the sensitivity of V_t to an infinitesimal bump of the time- t implied volatility $\sigma_k(t) := \sigma_{\text{imp}}(t, t + \tau_k)$ corresponding to an ATM swaption maturing at $t + \tau_k$, conditional on \mathcal{F}_t . This is graphically represented in figure 7(d).

4.3. Computing the sensitivity

First we consider the general case and denote

$$\theta_t := (\theta_1(t), \theta_2(t), \dots, \theta_K(t))^\top$$

where either $\theta_k(t) := R_k(t)$ or $\theta_k(t) := \sigma_k(t)$. Later on we will distinguish between the two cases. Without loss of generality, assume that $t \in (T_{m-1}, T_m]$. The aim is to compute

$$D_\theta \Pi(t) \Big|_{\theta=\theta_t} = \sum_{i=m}^{M-1} D_\theta \Pi_i(t) \Big|_{\theta=\theta_t}$$

We iteratively compute the sensitivities of each subportfolio, for $i = M-1, \dots, m$. Starting with $i = M-1$, a static replication of $V_{T_{M-1}}$ is achieved with a single European swaption, i.e. $\Pi_{M-1} := \pi_{M-1,1}$. Thus, setting $w_{M-1,1} = 1$ and $K_{M-1,1} = K$. The Delta and Vega simply correspond to the sensitivities of a European swaption, i.e.

$$D_\theta \Pi_{M-1}(t) = \left(\frac{\partial \pi_{M-1}(t)}{\partial \theta_1}, \dots, \frac{\partial \pi_{M-1}(t)}{\partial \theta_K} \right)$$

Next, consider the case $i = m < M-1$ and assume the sensitivities $D_\theta \Pi_{m+1}(t), \dots, D_\theta \Pi_{M-1}(t)$ are known. With slight abuse of notation, we have according to the chain rule the following relation

$$D_\theta \Pi_m(t) = \underbrace{\frac{\partial \Pi_m(t)}{\partial \theta}}_{\text{options Delta/Vega}} + \underbrace{\frac{\partial \Pi_m(t)}{\partial \mathbf{w}_m} \frac{\partial \mathbf{w}_m}{\partial \theta}}_{\text{weight sensitivities}} + \underbrace{\frac{\partial \Pi_m(t)}{\partial \mathbf{k}_m} \frac{\partial \mathbf{k}_m}{\partial \theta}}_{\text{strike sensitivities}} \quad (18)$$

We detail each term below:

Options Delta The first term represents the standard Delta or Vega of the options composing the portfolio, i.e.

$$\frac{\partial \Pi_m(t)}{\partial \theta} = \mathbf{w}_m^\top D_\theta \boldsymbol{\pi}_m(t) \quad (19)$$

Weight sensitivities The weights \mathbf{w}_m are optimized, such that the portfolio payoff matches a target function (equation (8)). This target function is sensitive to parameter changes in the zero-rate curve

or volatility term-structure. As a consequence we need to consider the weight sensitivities, i.e.

$$\frac{\partial \Pi_m(t)}{\partial \mathbf{w}_m} \frac{\partial \mathbf{w}_m}{\partial \theta} = \boldsymbol{\pi}_m^\top(t) D_\theta \mathbf{w}_m \quad (20)$$

Strike sensitivities For the same reason as for the weights, the sensitivities for the strikes should be considered. However, we conjecture the impact of the strike sensitivities is relatively small, which is supported by our numerical experiments. We will ignore this term for the remainder of this work.

A further derivation of the Delta/Vega sensitivities is the subject of sections 4.4 and 4.5. A further derivation of the weight sensitivities is subject of section 4.6.

4.4. The delta sensitivities

We start by working out the terms of the form $D_\theta \boldsymbol{\pi}_m$ and consider the case $\theta_k(t) = R_k(t) := R(t, t + \tau_k)$. We perform the derivation for the general setting $0 \leq t \leq T \leq T_m$ and aim to compute

$$D_\theta \boldsymbol{\pi}_m(T) \Big|_{\theta=\theta_t} = \begin{pmatrix} \frac{\partial \pi_{m,1}(T)}{\partial R_1(t)} & \cdots & \frac{\partial \pi_{m,1}(T)}{\partial R_K(t)} \\ \vdots & \ddots & \vdots \\ \frac{\partial \pi_{m,n}(T)}{\partial R_1(t)} & \cdots & \frac{\partial \pi_{m,n}(T)}{\partial R_K(t)} \end{pmatrix}$$

The $D_\theta \boldsymbol{\pi}_m$ in (19) corresponds to the case $t = T$. The case $t < T$ will serve as a path-wise sensitivity estimator, which will be required to compute the weight sensitivities in the subsequent section.

We first characterize the forward Delta sensitivity of a zero-coupon bond in the following lemma.

LEMMA 4.1 *Let $0 \leq t \leq T_0 \leq T_1$ and assume that the zero-rate curve $\tau \mapsto R(t, t + \tau)$ is an interpolation between nodes $\{R_0(t), \dots, R_K(t)\}$, where $R_k(t) := R(t, t + \tau_k)$. Then the forward Delta sensitivity is given by*

$$\frac{\partial P(T_0, T_1)}{\partial R_k(t)} = \alpha_k(\tau_0, \tau_1) P(T_0, T_1)$$

where $\tau_i = T_i - t$ and α_k denotes an \mathcal{F}_t -measurable coefficient, dependent on the interpolation scheme of the zero-rate nodes. If the interpolation is linear, the coefficient is given by

$$\alpha_k(\tau_0, \tau_1) := \alpha'_k(\tau_1) - \alpha'_k(\tau_0),$$

$$\alpha'_k(\tau) := \begin{cases} \frac{\tau - \tau_{k-1}}{\tau_k - \tau_{k-1}} \tau & \text{if } \tau \in (\tau_{k-1}, \tau_k] \\ \frac{\tau_{k+1} - \tau}{\tau_{k+1} - \tau_k} \tau & \text{if } \tau \in (\tau_k, \tau_{k+1}] \\ 0 & \text{otherwise} \end{cases}$$

A proof to this lemma is given in appendix 1. Now, for brevity denote $H_{m,i}(t) := \mathbb{E}^{m,M}[(\delta_{m,i}(S_{m,M}(T_m) - K_{m,i}))^+ | \mathcal{F}_t]$.

Then by an application of the chain rule and the product rule, we can write

$$\begin{aligned} \frac{\partial \pi_{m,i}(T)}{\partial R_k(t)} &= A_{m,M}(T) \underbrace{\frac{\partial H_{m,i}(T)}{\partial S_{m,M}(T)}}_{\text{Bachelier Delta}} \underbrace{\frac{\partial S_{m,M}(T)}{\partial R_k(t)}}_{\text{Swaprte sensitivity}} \\ &\quad + H_{m,i}(T) \underbrace{\frac{\partial A_{m,M}(T)}{\partial R_k(t)}}_{\text{Annuity sensitivity}} \end{aligned} \quad (21)$$

Working out each of the three partial derivatives above, we find

Bachelier Delta Under the swap rate dynamics approximation given in (15), this term is computed as the standard Bachelier Delta sensitivity, i.e.

$$\begin{aligned} \frac{\partial H_{m,i}(T)}{\partial S_{m,M}(T)} &= \delta_{m,i} \Phi(\delta_{m,i} d), \\ d &= \frac{S_{m,M}(T) - K_{m,i}}{\sigma_{\text{imp}}(T, T_m) \sqrt{T_m - T}} \end{aligned} \quad (22)$$

Annuity sensitivity From Lemma 4.1 it directly follows

$$\frac{\partial A_{m,M}(T)}{\partial R_k(t)} = \sum_{j=m+1}^M \tau_j \alpha_k(T - t, T_j - t) P(T, T_j)$$

Swap rate sensitivity Applying the chain rule and Lemma 4.1, it follows

$$\begin{aligned} \frac{\partial S_{m,M}(T)}{\partial R_k(t)} &= \frac{\alpha_k(T - t, T_m - t) P(T, T_m) - \alpha_k(T - t, T_m - t) \times P(T, T_m)}{A_{m,M}(T)} \\ &\quad - S_{m,M}(T) \frac{\sum_{j=m+1}^M \tau_j \alpha_k(T - t, T_j - t) P(T, T_j)}{A_{m,M}(T)} \end{aligned}$$

4.5. The vega sensitivities

Now consider the case $\theta_k(t) = \sigma_k(t) := \sigma_{\text{imp}}(t, t + \tau_k)$. Just as for Delta, we perform the derivation for the general setting $0 \leq t \leq T \leq T_m$ and aim to compute

$$D_{\theta} \pi_m(T) \Big|_{\theta=\theta_t} = \begin{pmatrix} \frac{\partial \pi_{m,1}(T)}{\partial \sigma_1(t)} & \cdots & \frac{\partial \pi_{m,1}(T)}{\partial \sigma_k(t)} \\ \vdots & \ddots & \vdots \\ \frac{\partial \pi_{m,n}(T)}{\partial \sigma_1(t)} & \cdots & \frac{\partial \pi_{m,n}(T)}{\partial \sigma_k(t)} \end{pmatrix}$$

Again, the $D_{\theta} \pi_m$ in (19) corresponds to the case $t = T$. The case $t < T$ will serve as a path-wise sensitivity estimator, which will be required to compute the weight sensitivities in the subsequent section.

We first characterize the forward implied volatility Vega in the following two lemmas.

LEMMA 4.2 Let $S_{m,M}$ satisfy the Gaussian approximation of (15) and denote $\tilde{\sigma}(t, T) := \sqrt{\frac{1}{T-t} \int_t^T \|\Sigma_{m,M}^{\top}(u) \eta(u)\|_2^2 du}$ for $0 \leq t \leq T \leq T_m$. Then

$$\begin{aligned} \frac{\partial \sigma_{\text{imp}}(T, T_m)}{\partial \sigma_k(t)} &= \frac{(T_m - t) \tilde{\sigma}(t, T_m) \frac{\partial \tilde{\sigma}(t, T_m)}{\partial \sigma_k(t)} - (T - t) \tilde{\sigma}(t, T) \frac{\partial \tilde{\sigma}(t, T)}{\partial \sigma_k(t)}}{(T_m - T) \sigma_{\text{imp}}(T, T_m)} \end{aligned}$$

A proof is given in appendix 2. The sensitivity of $\tilde{\sigma}(t, T)$ w.r.t. $\sigma_k(t)$ is an \mathcal{F}_t -measurable random variable, which is dependent on the calibration procedure of the instantaneous volatility $\eta(t)$. If η_1, \dots, η_K are scalars obtained by a bootstrap procedure, the sensitivities can be further characterized by the following lemma.

LEMMA 4.3 Let $\eta(t)$ be scalar-valued and piece-wise constant. Let $\Sigma_{m,M}$ be as in (14) and let Σ_k denote the diffusion of the swap rate underlying the option corresponding to $\sigma_k(t) := \sigma_{\text{imp}}(t, t + \tau_k)$. Then the forward implied volatility Vega is given by

$$\frac{\partial \tilde{\sigma}(t, T)}{\partial \sigma_k(t)} = \alpha_k(T) \frac{\tau_k \sigma_k(t)}{\tau \tilde{\sigma}(t, T)}$$

where $\tau := T - t$ and α_k denotes an \mathcal{F}_t -measurable coefficient. If by approximation $\Sigma_{m,M} \approx \Sigma_k \approx \Sigma_{k+1}$, the coefficient is given by

$$\alpha_k(T) = \begin{cases} \frac{\int_{T_{k-1}}^T \Sigma_{m,M}^2(u) du}{\int_{T_{k-1}}^{T_k} \Sigma_k^2(u) du} & \text{if } T_{k-1} < T \leq T_k \\ \frac{\int_T^{T_{k+1}} \Sigma_{m,M}^2(u) du}{\int_{T_k}^{T_{k+1}} \Sigma_{k+1}^2(u) du} & \text{if } T_k < T \leq T_{k+1} \\ 0 & \text{otherwise} \end{cases}$$

where we denote $T_k := t + \tau_k$.

A proof is given in appendix 3. By an application of the chain rule, we obtain

$$\begin{aligned} \frac{\partial \pi_{m,i}(T)}{\partial \sigma_k(t)} &= \underbrace{\frac{\partial \pi_{m,i}(T)}{\partial \sigma_{\text{imp}}(T, T_m)}}_{\text{Bachelier Vega}} \underbrace{\frac{\partial \sigma_{\text{imp}}(T, T_m)}{\partial \sigma_k(t)}}_{\text{Imp.vol. sensitivity}} \\ &\quad + \underbrace{\frac{\partial \pi_{m,i}(T)}{\partial S_{m,M}(T)}}_{\text{Bachelier Delta}} \underbrace{\frac{\partial S_{m,M}(T)}{\partial \sigma_k(t)}}_{\text{Swaprte sensitivity}} \end{aligned} \quad (23)$$

Considering each of the partial derivatives above, we find **Bachelier Vega** Under the swap rate approximation given in (15), this term can be computed as the Bachelier Vega sensitivity, i.e.

$$\begin{aligned} \frac{\partial \pi_{m,i}(T)}{\partial \sigma_{\text{imp}}(T, T_m)} &= A_{m,M}(T) \sqrt{T_m - T} \varphi(d), \\ d &= \frac{S_{m,M}(T) - K_{m,i}}{\sigma_{\text{imp}}(T, T_m) \sqrt{T_m - T}} \end{aligned}$$

Bachelier Delta Under the swap rate approximation given in (15), this term can be computed as the Bachelier Delta sensitivity. Hence, similar to (22), we have

$$\begin{aligned} \frac{\partial \pi_{m,i}(T)}{\partial S_{m,M}(T)} &= \delta_{m,i} A_{m,M}(T) \Phi(\delta_{m,i} d), \\ d &= \frac{S_{m,M}(T) - K_{m,i}}{\sigma_{\text{imp}}(T, T_m) \sqrt{T_m - T}} \end{aligned} \quad (24)$$

Imp. vol. sensitivity The forward implied volatility sensitivity is characterized in Lemma 4.2. Under the assumptions of Lemma 4.3, the sensitivity can be approximated as

$$\frac{\partial \sigma_{\text{imp}}(T, T_m)}{\partial \hat{\sigma}_k(t)} = \frac{(\alpha_k(T_m) - \alpha_k(T)) \tau_k \sigma_k(t)}{(T_m - T) \sigma_{\text{imp}}(T, T_m)}$$

Swap rate sensitivity Under the swap rate approximation given in (15), we can write

$$S_{m,M}(T) \sim S_{m,M}(t) + \tilde{\sigma}(t, T) (W_T^{m,M} - W_t^{m,M})$$

where $\tilde{\sigma}(t, T) := \sqrt{\frac{1}{T-t} \int_t^T \|\Sigma_{m,M}^\top(u) \eta(u)\|_2^2 du}$. It follows that the forward swap rate sensitivity can be estimated as

$$\frac{\partial S_{m,M}(T)}{\partial \sigma_k(t)} = \frac{S_{m,M}(T) - S_{m,M}(t)}{\tilde{\sigma}(t, T)} \frac{\partial \tilde{\sigma}(t, T)}{\partial \sigma_k(t)}$$

Under the assumptions of Lemma 4.3, this can be further approximated as

$$\frac{\partial S_{m,M}(T)}{\partial \sigma_k(t)} = \frac{\tau_k \alpha_k(T)}{T - t} (S_{m,M}(T) - S_{m,M}(t))$$

4.6. The weight sensitivities

The weights and strikes of Π_m are obtained through an optimization scheme as described in section 3.1. The objective function for this optimization can equivalently be expressed as

$$L_m(\mathbf{w}, \mathbf{k}) = \mathbb{E}^\mathbb{Q} \left[\left(\mathbf{w}^\top \boldsymbol{\pi}_m(T_m) - g_m(T_m) \right)^2 \right]$$

where we denoted $g_m(t) := \frac{f_m(x_t)}{q_m(x_t)}$. The parameters \mathbf{w}_m and \mathbf{k}_m are determined by

$$(\mathbf{w}_m, \mathbf{k}_m) = \arg \min_{\mathbf{w}, \mathbf{k}} \left\{ L_m(\mathbf{w}, \mathbf{k}) \mid \mathbf{w}, \mathbf{k} \in \mathbb{R}^n \right\} \quad (25)$$

Assuming L_m attains a local minimum in $(\mathbf{w}_m, \mathbf{k}_m)$, it follows by the first order optimality condition that $\nabla L_m = \mathbf{0}$. In particular, we have

$$\lambda_m(\mathbf{w}_m) := \left(\frac{\partial L_m}{\partial w_1}, \dots, \frac{\partial L_m}{\partial w_n} \right) \Big|_{\mathbf{w}=\mathbf{w}_m} = \mathbf{0}$$

where we can compute

$$\begin{aligned} \lambda_m(\mathbf{w}) &= \mathbb{E}^\mathbb{Q} \left[D_w \left(\mathbf{w}^\top \boldsymbol{\pi}_m(T_m) - g_m(T_m) \right)^2 \right] \\ &= 2\mathbb{E}^\mathbb{Q} \left[\left(\mathbf{w}^\top \boldsymbol{\pi}_m(T_m) - g_m(T_m) \right) \boldsymbol{\pi}_m(T_m) \right] \end{aligned}$$

In order to characterize the weight sensitivities, we first consider the Jacobian of λ_m (i.e. the Hessian of L_m w.r.t. \mathbf{w}). The Jacobian can be obtained by again differentiating the expression above, which yields

$$D_w \lambda_m = 2\mathbb{E}^\mathbb{Q} \left[D_w \left(\mathbf{w}^\top \boldsymbol{\pi}_m(T_m) - g_m(T_m) \right) \boldsymbol{\pi}_m(T_m) \right] \quad (26)$$

$$= 2\mathbb{E}^\mathbb{Q} \left[\boldsymbol{\pi}_m(T_m) \boldsymbol{\pi}_m^\top(T_m) \right] \quad (27)$$

Hence, note that $D_w \lambda_m$ is specified by the auto-correlation matrix of the random vector $\boldsymbol{\pi}_m$. In other words, it depends on the cross-sectional co-variance matrix of the portfolio pay-offs. With this expression at hand, we characterize the weight sensitivities in the following proposition

PROPOSITION 4.4 *Consider a swaption portfolio Π_m written on swap rate $S_{m,M}$ and let its weights and strikes satisfy equation (25). Furthermore assume that the strikes K_1, \dots, K_n are all distinct. Then the derivative of \mathbf{w}_m w.r.t. $\boldsymbol{\theta}_t$ is well-defined and its Jacobian matrix at $(\mathbf{w}_m, \boldsymbol{\theta}_t)$ is given by*

$$\frac{\partial \mathbf{w}_m}{\partial \boldsymbol{\theta}_t} = - (D_w \lambda_m)^{-1} D_\theta \lambda_m \Big|_{\mathbf{w}=\mathbf{w}_m, \boldsymbol{\theta}=\boldsymbol{\theta}_t} \quad (28)$$

This result is a consequence of the *implicit function theorem*. A rigorous proof is given in appendix 4.

REMARK 1 The first term in (28) is given by the inverse of the matrix in (27). Note that this term is both path-independent and time-independent. Hence, although its computation requires numerical estimation and matrix inversion, it can be used along each MC path, at each time-step without nested re-evaluation.

What remains to be done is computing $D_\theta \lambda_m$. We therefore compute

$$\begin{aligned} D_\theta \lambda_m &= 2\mathbb{E}^\mathbb{Q} \left[\boldsymbol{\pi}_m D_\theta \left(\mathbf{w}^\top \boldsymbol{\pi}_m(T_m) - g_m(T_m) \right) \right] \\ &\quad + 2\mathbb{E}^\mathbb{Q} \left[\left(\mathbf{w}^\top \boldsymbol{\pi}_m(T_m) - g_m(T_m) \right) D_\theta \boldsymbol{\pi}_m(T_m) \right] \end{aligned}$$

Note that the second term is zero if subportfolio Π_m yields a perfect fit to the target function g_m . Therefore, we can expect this term to be small and we will ignore it. Working out the first term, we find

$$D_\theta \lambda_m \approx 2\mathbb{E}^\mathbb{Q} \left[\boldsymbol{\pi}_m \mathbf{w}^\top D_\theta \boldsymbol{\pi}_m(T_m) \right] - 2\mathbb{E}^\mathbb{Q} \left[\boldsymbol{\pi}_m D_\theta g_m(T_m) \right]$$

In the first term of the expression above, $D_\theta \boldsymbol{\pi}_m(T_m)$ represent the Deltas/Vegas of subportfolio Π_m . This term needs to be numerically estimated. The path-wise estimator has been derived in sections 4.4 and 4.5. The second term also needs to be numerically estimated. The path-wise estimator involves $D_\theta g_m(T_m)$, where $g_m(t) := \frac{f_m(x_t)}{q_m(x_t)} = \max\{h_m(x_t) - \sum_{i=m+1}^{M-1} \Pi_i(t), 0\}$. Working out this term and denoting the event $E_m := \{h_m(x_{T_m}) > \sum_{i=m+1}^{M-1} \Pi_i(T_m)\}$ yields

$$D_\theta g_m(T_m) = \mathbb{1}_{E_m} D_\theta h_m(x_{T_m}) - \mathbb{1}_{E_m} \sum_{i=m+1}^{M-1} D_\theta \Pi_i(T_m) \quad (29)$$

In the expression above, the term $D_\theta h_m(x_{T_m})$ is equivalent to $D_\theta \pi_{m,0}(T_m)$ with $K_{m,0} = K$. The terms $D_\theta \Pi_i(T_m)$ for $i = m+1, \dots, M-1$ were assumed to be known from previous iterations.

Algorithm 1: The static replication algorithm

Generate N risk-factor scenarios for $x_{T_0}, \dots, x_{T_{M-1}}$;
for $n = 1, \dots, N$ **do**
 | $C_{M-1}(n) \leftarrow 0$;
end
Initialize parameters β_{M-1} from $\mathcal{U}(-1, 1)$;
for $m = M - 1, \dots, 1$ **do**
 for $n = 1, \dots, N$ **do**
 | $f_m(n) \leftarrow$
 | $q_m(x_{T_m}(n)) \cdot \max\{h_m(x_{T_m}(n)) - C_m(n), 0\}$;
 | $z_m(n) \leftarrow S_{m,M}(T_m, x_{T_m}(n))$;
 end
 $\beta_m \leftarrow \arg \min_{\beta} \frac{1}{N} \sum_{n=1}^N (G_m(z_m(n); \beta) - f_m(n))^2$;
 for $n = 1, \dots, N$ **do**
 | $C_{m-1}(n) \leftarrow$
 | $\sum_{j=m}^{M-1} M_{T_{m-1}} \mathbb{E}^{\mathbb{Q}} \left[\frac{G_j(z_j; \beta_j)}{M_{T_j, q_j}(x_{T_j})} \middle| x_{T_{m-1}}(n) \right]$;
 end
 Initialize $\beta_{m-1} \leftarrow \beta_m$;
end
Result: $V(0) = \mathbb{E}^{\mathbb{Q}} \left[\frac{C_0}{M_{T_0}} \middle| x_0 \right]$

4.7. The algorithms

Below we summarize the algorithms for composing the replicating portfolios (Algorithm 1) and computing the path-wise sensitivities (Algorithm 2). For the execution of the path-wise sensitivity algorithm, it is assumed that the static replication has been performed and subportfolios Π_0, \dots, Π_{M-1} have been derived from the trained parameters $\beta_0, \dots, \beta_{M-1}$.

5. Numerical experiments

We demonstrate the performance of the algorithms presented in this work by considering several case studies. We select a one-factor Gaussian model with piece-wise constant instantaneous volatility to perform our numerical experiments (see Andersen and Piterburg 2010). That is, we consider state-variable dynamics given by

$$\begin{aligned} dx_t &= (\phi(t) - ax_t) dt + \eta(t) dW_t, \quad x_0 = 0 \\ \phi(t) &= \int_0^t e^{-2a(t-u)} \eta(u) du \\ r_t &= f(0, t) + x_t \end{aligned}$$

where the scalar a denotes a constant mean-reversion rate, W_t is a one-dimensional Brownian motion under \mathbb{Q} and $f(t, T) := -\frac{\partial \log R(t, T)}{\partial T}$ denotes the instantaneous forward rate. The time-zero yield curve is assumed to be a linear interpolation between zero-rates with tenors τ_k as given in table 1. We consider a flat yield curve with all zero-rates equal to 3%. The instantaneous forward rate is assumed to be piece-wise constant between the tenor dates τ_k . Each volatility level η_k is bootstrapped against an ATM European payer swaption with maturity τ_k and implied volatility σ_k , such that $\tilde{\sigma}_{\text{imp}}(0, \tau_k) =$

Algorithm 2: The path-wise sensitivity algorithm

Generate N risk-factor scenarios for x_{t_0}, \dots, x_{t_J} ;
for $m = M - 1, \dots, 1$ **do**
 | $A_{1,m} \leftarrow \left(\mathbb{E}^{\mathbb{Q}} [\boldsymbol{\pi}_m(T_m) \boldsymbol{\pi}_m^{\top}(T_m)] \right)^{-1}$;
end
for $t_j = 0, t_1, \dots, t_J$ **do**
 $m_j \leftarrow \min\{0, \dots, M - 1 : t_j \leq T_m\}$;
 for $m = M - 1, \dots, m_j$ **do**
 | $\theta \leftarrow (R_1(t_j), \dots, R_K(t_j))^{\top}$ OR
 | $(\sigma_1(t_j), \dots, \sigma_K(t_j))^{\top}$;
 | $A_{2,m} \leftarrow \mathbb{E}^{\mathbb{Q}} [\boldsymbol{\pi}_m \boldsymbol{w}^{\top} D_{\theta} \boldsymbol{\pi}_m(T_m)]$, using
 | equation (21) / (23);
 | $A_{3,m} \leftarrow \mathbb{E}^{\mathbb{Q}} [\boldsymbol{\pi}_m D_{\theta} g_m(T_m)]$, using
 | equation (29) and (21) / (23);
 | $D_{\theta} \boldsymbol{w}_m \leftarrow -A_{1,m} (A_{2,m} - A_{3,m})$;
 for $n = 1, \dots, N$ **do**
 | $D_{\theta} \boldsymbol{\pi}_m(n) \leftarrow D_{\theta} \boldsymbol{\pi}_m(t_j; x_{t_j}(n))$ using
 | equation (21) / (23);
 | $D_{\theta} \Pi_m(n) \leftarrow$
 | $\boldsymbol{w}_m^{\top} D_{\theta} \boldsymbol{\pi}_m(n) + \boldsymbol{\pi}_m^{\top}(t_j; x_{t_j}(n)) D_{\theta} \boldsymbol{w}_m$;
 end
 end
 for $n = 1, \dots, N$ **do**
 | $D_{\theta} \Pi(t_j, n) \leftarrow \sum_{m=m_j}^{M-1} D_{\theta} \Pi_m(n)$;
 end
end
Result: $\{D_{\theta} \Pi(t_j, n) : j = 0, \dots, J; n = 1, \dots, N\}$

σ_k . All implied volatility targets σ_k are set to 50 bps, resulting in the instantaneous volatilities presented in table 1. The mean-reversion rate is fixed a priori at $a = 0.01$. The tenors τ_k have been selected to coincide with the SIMM tenor buckets specified by ISDA (International Swaps and Derivatives Association, Inc 2020).

Sampling required for regression data, pricing, sensitivities, exposures and IM profiles is done using Monte Carlo simulation. The risk factor is simulated through an Euler discretization scheme. The discretization uses weekly time-steps, i.e. $\Delta t = 1/52$.

As case studies, we consider a $1Y \times 5Y$ and a $1Y \times 10Y$ Bermudan receiver swaption with unit notional. The underlying of the $T_0 \times T_M$ Bermudan is assumed to be a fixed-for-floating interest rate swap with fixed rate K and a tenor structure $\mathcal{T} = \{T_0, \dots, T_M\}$ of annual payments. Unless stated otherwise, the fixed rate K is selected to be ATM, i.e. $K = K_{\text{ATM}} := S_{0,M}(0)$. The exercise dates of the Bermudan are set to coincide with the fixing dates of the underlying swap, i.e. T_0 being the first and T_{M-1} being the last exercise opportunity.

The composition of the static replication portfolio relies on the regression of feed-forward neural networks with a single hidden layer. Below, we list some details concerning the fitting procedure of the neural networks.

- As an independent regression variable for G_m we select the swap rate $S_{m,M}(T_m)$. The regression data

Table 1. Model parameters.

τ_k	2W	1M	3M	6M	1Y	2Y	3Y	5Y	10Y	15Y	20Y	30Y
R_k	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
η_k (bps)	50.9	50.9	50.9	51.1	51.2	51.2	51.2	51.3	51.3	57.2	57.8	61.0

is normalized before G_m is fitted to the data, to improve the regression robustness (see Hoencamp *et al.* 2023 for details).

- Each network G_m is selected to have 8 hidden nodes.
- Each regression is performed with a data set of 2000 independently sampled training points.
- The neural network parameters are optimized using the AdaMax optimizer, which is a stochastic gradient-descent method (Kingma and Ba 2014). The batch-size (i.e. number of training points used per iteration) is set to 32. The learning rate (i.e. the step-size scaling per iteration) is set to 0.01.
- The parameters of G_{M-1} are randomly initialized from a uniform distribution $\mathcal{U}(-1, 1)$. For $m < M - 1$, the parameters of G_m are initialized with the optimized parameters of G_{m+1} .

5.1. Pricing

We start by analyzing the pricing accuracy of the static replication. As each neural network has 8 hidden nodes, every

Table 2. V_0 estimates Bermudan swaption for different levels of moneyness, presented in basis points of the notional.

$K/S_{0,M}$	$1Y \times 5Y$		$1Y \times 10Y$	
	SR	LSM (SE)	SR	LSM (SE)
60%	14.43	14.18 (0.15)	61.04	60.20 (0.48)
80%	46.63	46.63 (0.35)	140.7	139.9 (0.8)
100%	130.4	130.3 (0.5)	300.5	300.2 (1.2)
120%	300.4	300.6 (0.7)	590.6	590.2 (1.4)
140%	543.8	544.6 (0.9)	1018.3	1018.8 (1.6)

regression yields a subportfolio of 8 European swaptions. Hence at t -zero, the $1Y \times 5Y$ Bermudan is replicated with 40 swaptions and the $1Y \times 10Y$ Bermudan is replicated with 80 swaptions.

5.1.1. Time-zero prices. Table 2 reports the V_0 price estimates implied by the risk-neutral valuation of the static replication portfolio (SR) for different levels of moneyness. As a benchmark, we report price estimates obtained from the least-square method (LSM) introduced by Longstaff and Schwartz (2001). As basis functions for the ordinary least-square regression, we use $\{1, x_t, x_t^2, x_t^3\}$. The LSM estimate is the mean of 25 independent runs of 80 000 MC paths each. In brackets, we report the standard error generated by these 25 runs.

From table 2, we observe a close correspondence between SR and LSM, in particular for ATM strikes. The accuracy slightly deteriorates for far OTM options. This can be attributed to a limited number of ITM paths at future monitor dates, yielding only a few non-zero data-points for the regression. SR is not subject to a standard error as the portfolios are evaluated at t -zero. SR is however subject to numerical errors due to the inaccuracy of the neural network regression. On the normalized data, the fits typically show a mean squared error in the order of $10^{-6} - 10^{-5}$.

5.1.2. Exposure profiles. Consecutively, we consider the exposure profile and related CVA quantities of the Bermudan swaptions. Exposure depends on the future distribution of the option price. Figure 8 shows the expected positive exposure (EPE) simulated with the static replication next to an LSM implied benchmark. The benchmark is generated with the *LSM-bundle* algorithm, presented in Feng *et al.* (2016). We briefly summarize this method below.

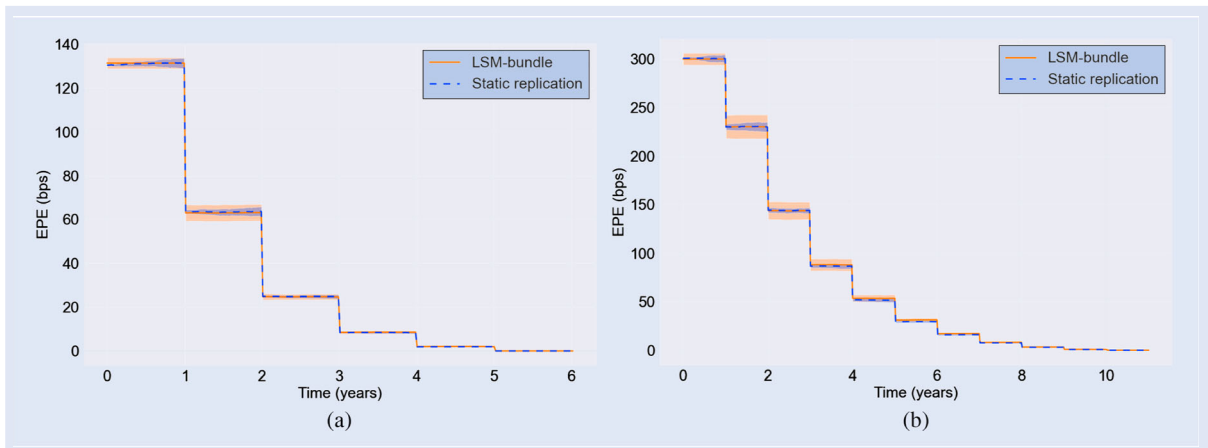


Figure 8. Expected positive exposure profile: LSM-bundle (orange) vs. static replication (blue). The LSM and static replication standard error windows are represented by the orange and blue shaded area, respectively, based on 25 independent runs of 3000 MC paths. (a) $1Y \times 5Y$ Bermudan and (b) $1Y \times 10Y$ Bermudan.

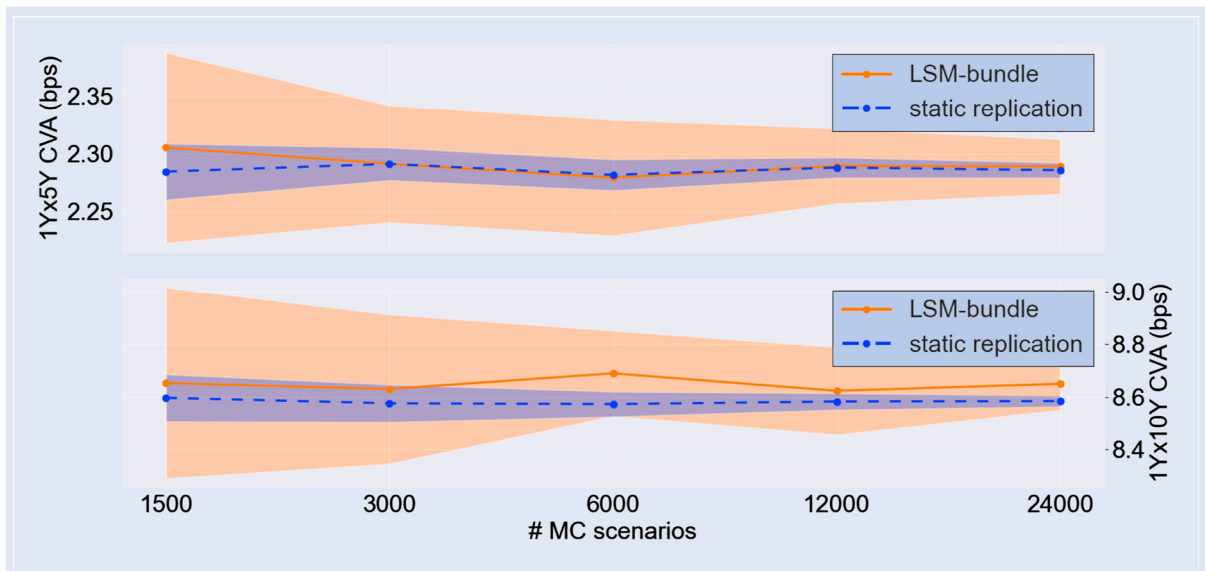


Figure 9. Convergence CVA for a $1Y \times 5Y$ (top) and $1Y \times 10Y$ (bottom) Bermudan swaption: LSM-bundle (orange) vs. static replication (blue). The LSM and static replication standard error windows are represented by the orange and blue shaded area, respectively.

The LSM-bundle algorithm allows one to estimate the EPE of a Bermudan option when combined with the LSM algorithm of Longstaff and Schwartz (2001). Within LSM, the continuation value $C_m(T_m)$ is estimated by regressing realized (discounted) cashflows at T_{m+1} against the risk-factor realizations at T_m on in-the-money MC paths. However, to compute the exposure, an unbiased estimate of the continuation value is required on *each* MC path on every exposure date (not just monitor dates). In Feng *et al.* (2016) it is therefore proposed to first perform an LSM sweep. Then in a second sweep, for each exposure date t the MC paths are divided into two bundles. If $t \in (T_{m-1}, T_m]$, the two bundles are defined as:

$$\mathcal{U}_1(t) := \{x_t : h_m(x_t) > 0\}, \quad \mathcal{U}_2(t) := \{x_t : h_m(x_t) \leq 0\}$$

Subsequently for both bundles an ordinary least-square regression is performed against the discounted realized cashflows that were computed in the first LSM sweep. The regression functions provide an unbiased estimate of the exposure for a given realization of the risk-factor in their respective domains. We refer to Feng *et al.* (2016) for details. Note that for the SR method, no second sweep of regressions is required. Exposures are obtained by simply pricing the replicating portfolio along the MC path.

The EPE profiles in figure 8 represent the mean of 25 MC simulations of 3000 paths each. The shaded areas represent a window of one standard error for both respective methods. We observe the typical staircase nature that corresponds to the exposure profile of a Bermudan option. Each discontinuity represents the passing of an exercise opportunity, after which the option is exercised for a certain fraction of the scenarios. Similar to the time-zero prices, we observe a good correspondence between the SR estimates and the benchmark. Additionally, it is noted that the standard error for SR is significantly smaller compared to LSM-bundle. This can be attributed to the benefit of the regress-later approach of SR compared to the regress-now approach of LSM. In the latter, the continuation value is estimated by regressing the

(noisy) realizations of the cashflow. In the former, the regression is performed against the deterministic option's target value (numerical errors aside). As a result, the MC noise is integrated out more efficiently with SR. In general, this indicates that a significantly smaller number of MC paths would be required to reach a similar level of accuracy.

We further study the convergence of the standard error, by considering the CVA statistic. To do so, we assume a constant $LGD = 1$ and constant hazard rate of $\lambda(t) = 0.01$. The CVA is estimated using equation (3). Figure 9 shows the mean and standard deviation of the CVA estimate generated with 25 independent MC runs, as a function of the number of MC paths per run. We observe that on average the SR standard error is a factor 4 smaller compared to the LSM-bundle benchmark. For an increasing number of paths, the two methods show a consistent convergence.

5.2. Sensitivities

We continue by analyzing the accuracy in the sensitivity computation for the static replication. We distinguish between time-zero sensitivities and future sensitivities along the MC path. The Delta and Vega sensitivities in this section are understood in accordance with the definitions provided in section 4.2.

5.2.1. Time-zero sensitivities. Figure 10 reports the time-zero Delta sensitivities for the 12 tenor buckets. The sensitivities have been scaled to represent the basis point value w.r.t. the zero-rate R_k . The SR results (presented in blue) are calculated in accordance with the Delta routine described in section 4.2. As a benchmark, we consider an LSM bump-and-reval estimate (presented in orange). For Delta, this means a zero-rate corresponding to one of the predefined tenors is bumped with one basis point and the yield curve is re-interpolated accordingly. The same seed is used for the LSM base run as for the bumped LSM run, to facilitate variance reduction.

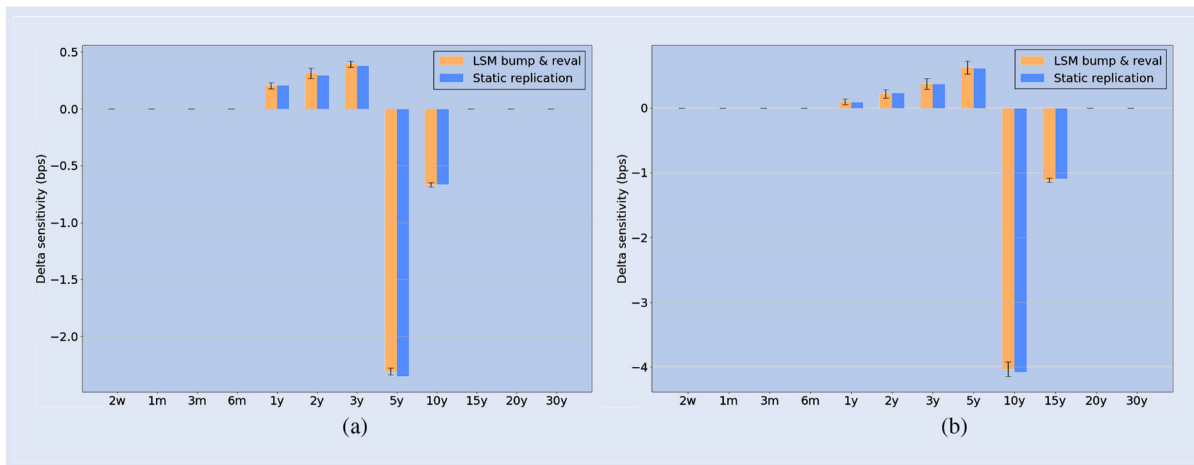


Figure 10. Bucketed Delta sensitivities: LSM bump-and-reval (orange) vs. static replication (blue). The LSM benchmark is based on 25 independent runs of 80,000 MC paths. (a) 1Y × 5Y Bermudan and (b) 1Y × 10Y Bermudan.




Figure 11. Bucketed Vega sensitivities: LSM bump-and-reval (orange) vs. static replication (blue). The LSM benchmark is based on 80,000 MC paths. (a) 1Y × 5Y Bermudan and (b) 1Y × 10Y Bermudan.

Figure 11 reports the time-zero Vega sensitivities for the 12 tenor buckets. These sensitivities have been scaled to represent the basis point value w.r.t. implied volatility σ_k . The bump-and-reval benchmark for Vega is obtained by bumping one of the implied volatilities with one basis point and re-bootstrapping the volatility term-structure accordingly. Also here the same seed is used for the base and the bumped LSM run.

The results shown in figures 10 and 11 are the mean of 25 independent runs of 80,000 paths each. For the LSM bump-and-reval, this means 13 LSM simulations were involved for each run (1 base simulation and 12 bumped simulations). The SR estimate only requires a single simulation, in order to estimate the weight sensitivities of the replicating portfolio. The swaption sensitivities can be readily evaluated. The error bars indicate the standard error of the LSM estimate. The SR standard errors were each 2 to 3 orders of magnitude smaller compared to LSM and are therefore not included in the figures. Overall, we observe a satisfactory correspondence between the SR sensitivities and the LSM benchmark. For the LSM Vega, we observe relatively large standard errors compared to Delta. This can be attributed to the fact that a

bumped yield curve translates one-to-one into bumped zero-coupon bond prices. A bumped volatility term-structure only indirectly impacts the risk-factor distribution and is therefore more susceptible to a higher standard error.

Table 3 reports the SIMM Delta, Vega and Curvature margins computed with the sensitivity estimates of SR and LSM, respectively, for different levels of moneyness. In brackets we report the standard error of the LSM implied quantities. For the Delta and Vega margin estimates, we see the accuracy of the SR method confirmed. The Curvature margin LSM estimates demonstrate a significantly larger standard error in the relative sense. Curvature margin reflects the Gamma sensitivity of the option, which in SIMM is approximated through a Vega-Gamma relationship (International Swaps and Derivatives Association, Inc 2016). Second-order sensitivity approximations are notorious for amplifying numerical errors of the first order sensitivities, which we see reflected in the LSM benchmark. The SR estimate does not suffer from such standard errors.

The convergence of the margin estimates is demonstrated in figure 12. Here we present the mean and standard error as a function of the number of MC paths, obtained from 25

Table 3. IM estimate Bermudan swaption.

$K/S_{0,M}$	Margin	$1Y \times 5Y$		$1Y \times 10Y$	
		SR	LSM (SE)	SR	LSM (SE)
60%	Delta	17.44	17.09 (1.19)	53.45	51.70 (4.07)
	Vega	7.81	8.09 (0.43)	26.01	26.34 (1.84)
	Curv.	1.12	1.42 (0.84)	2.20	3.57 (3.22)
80%	Delta	49.07	47.51 (2.91)	112.1	107.0 (6.9)
	Vega	14.93	15.54 (0.62)	38.98	39.28 (1.79)
	Curv.	2.75	2.91 (1.63)	4.23	4.10 (3.62)
100%	Delta	115.2	111.9 (3.8)	213.4	211.3 (14.7)
	Vega	19.56	20.04 (1.00)	45.46	46.58 (3.34)
	Curv.	5.06	5.61 (2.8)	7.21	8.79 (6.42)
120%	Delta	200.5	198.3 (4.7)	355.7	355.5 (14.5)
	Vega	12.84	12.80 (1.13)	33.58	34.22 (3.55)
	Curv.	4.29	3.63 (2.96)	8.03	10.84 (8.99)
140%	Delta	247.3	246.7 (2.4)	464.0	463.1 (10.0)
	Vega	2.92	3.09 (0.74)	10.54	10.61 (2.57)
	Curv.	1.11	1.58 (1.55)	3.27	4.82 (5.37)

independent runs. We observe that the consistency of the SR method is confirmed for an increasing number of MC paths.

5.2.2. Future sensitivities: European swaption. We finally analyze the expected future sensitivity profiles of a Bermudan swaption. Calculation of expected future sensitivities through

the LSM bump-and-reval method would impose a severe computational burden. A nested MC simulation is required for each time-step, on each MC path, for each sensitivity. Benchmarking future sensitivity profiles is hence not feasible with LSM and out of scope for this paper.

To verify the accuracy of the SR method along the MC path, we therefore start by considering a $5Y \times 5Y$ European payer swaption. This can be considered a special case of the Bermudan, with only a single exercise opportunity at T_0 . We proceed to replicate this swaption, while selecting the $15Y$ swap rate as regression asset. As a result, the $5Y \times 5Y$ swaption is replicated with a portfolio of $5Y \times 15Y$ swaptions. The same neural network structure is used as for the Bermudans. Subsequently, we analyze the sensitivity profiles of the replicating portfolio, which we benchmark with the analytical sensitivities of the swaption.

The results are displayed in figure 13. The Vega profiles on the right show a near-perfect match between SR and the benchmark. We do observe some numerical errors in the Delta $10Y$, $15Y$ and $20Y$ buckets. Around $t = 0$, the latter two are dominant for the $5Y \times 15Y$ swaptions in the replicating portfolio, but zero for the target $5Y \times 5Y$ swaption. If the replication is perfect, the sensitivities in this bucket cancel with the weight sensitivities. Here we note that some small errors remain. Yet, overall, we observe a satisfactory agreement between SR and the analytical profiles.




Figure 12. Convergence Delta (top), Vega (middle) and Curvature (bottom) IM: LSM bump-and-reval (orange) vs. static replication (blue). The LSM standard error window is represented by the shaded orange area. (a) $1Y \times 5Y$ Bermudan and (b) $1Y \times 10Y$ Bermudan.




Figure 13. Bucketed Delta (left) and Vega (right) profiles for static replication of a European $5Y \times 5Y$ swaption. SR sensitivities (solid) are shown next to an exact benchmark (dashed). (a) Delta profiles and (b) Vega profiles.

5.2.3. Future sensitivities: Bermudan swaption. Figure 14 shows the expected future Delta sensitivity profiles of a Bermudan swaption. The profiles are generated by taking the mean of 25 independent MC runs of 3000 paths each. The shaded areas (where visible) represent a single standard error window. As the underlying is a receiver swap, we note that the Delta profiles are mostly negative. The tenor buckets surrounding the time until maturity of the underlying swap are dominant. As time passes, the time to maturity decreases and the dominant sensitivities shift from one bucket to the next, implying the triangular-shaped profiles. We see this for example reflected in the (red-colored) Delta-profile of the 5Y bucket in figure 14(a). In the first year, the sensitivity grows (in the absolute sense) to its maximum until the maturity of the underlying swap is exactly 5Y away. Then, the sensitivity of this bucket starts to drop and flows into the 3Y bucket. The discontinuous jumps signify the passing of an exercise opportunity. Despite the small number of paths, we observe a satisfactory low standard error.

Figure 15 shows the corresponding expected Vega profiles and their standard error window. The sensitivities are scaled by the implied volatility in basis points of each corresponding bucket, which is in line with the SIMM routines. The Vega sensitivities are bucketed according to their time to expiry. Hence, we observe triangular patterns resulting from Vegas

shifting from one bucket into the next as the time until the future monitor dates decreases. Also here, we observe that the shaded areas representing the standard errors are minor or even invisible in the figure, even though the number of paths is relatively small.

Figure 16 shows the expected initial margin profiles, obtained from the simulated sensitivity distributions. The shaded areas again denote a single standard error window. The figure shows the separate profiles corresponding the Delta-, Vega- and Curvature margin. The sum of these three yields the total expected initial margin as a function of time. From this plot, we observe that overall the Delta margin is dominant for the computation of IM and MVA. For Delta, we observe the typical staircase profile. Each discontinuity represents the passing of an exercise date and the subsequent decrease in the option's total Delta sensitivity. The volume of Vega scales with the square root of the times until expiry of the underlying options. Hence the smooth decay of the Vega margin over time, with similar jumps after each monitor date. The contribution of the curvature margin to the total IM and MVA is very limited. In terms of shape, the profile shows a similar nature as Delta.

We study the convergence of the margin estimates by considering MVA as a function of the number of MC paths. Figure 17 shows the Delta-, Vega- and Curvature components




Figure 14. Bucketed Delta profiles for static replication. The standard error windows are represented by the shaded area, based on 3000 MC paths. (a) $1Y \times 5Y$ Bermudan and (b) $1Y \times 10Y$ Bermudan.




Figure 15. Bucketed Vega profiles for static replication. The standard error windows are represented by the shaded area, based on 3000 MC paths. (a) $1Y \times 5Y$ Bermudan and (b) $1Y \times 10Y$ Bermudan.

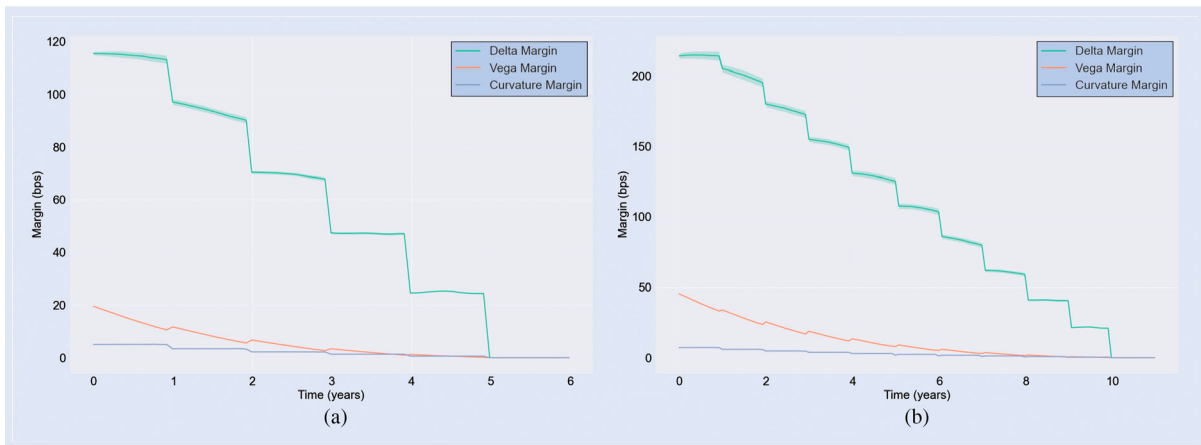


Figure 16. Delta (green), Vega (orange) and Curvature (blue) margin profiles for static replication. The standard error windows are represented by the shaded area, based on 3000 MC paths. (a) $1Y \times 5Y$ Bermudan and (b) $1Y \times 10Y$ Bermudan.




Figure 17. Convergence Delta (top), Vega (middle) and Curvature (bottom) MVA for static replication. The standard error windows are represented by the shaded areas. (a) $1Y \times 5Y$ Bermudan and (b) $1Y \times 10Y$ Bermudan.

of MVA in separate plots. We report the mean and standard error, obtained from 25 independent runs. The MVA estimates have been generated from the future IM distributions, while assuming a constant funding spread of $FS(t) = 0.01$. From the figures, we observe a consistent convergence and see the low standard errors of the sensitivity profiles confirmed. We conclude that generally, 3000 MC paths suffice to guarantee a relative standard error of below 1%.

6. Conclusion

This paper presented a static replication algorithm for Bermudan swaptions under an affine term-structure model. We showed that under the appropriate conditions, one can construct a portfolio of European swaptions that mirrors the value of the exotic instrument until it is either exercised or matured. This replication is static, i.e. despite the early-exercise features of the instrument, the portfolio does not need to be re-balanced throughout the lifetime of the original contract. We continued to derive efficient estimators for the price and sensitivities of the portfolio along the Monte Carlo path. These quantities are essential for the quantification of modern risk metrics such as exposure, CVA, IM and MVA.

Exploiting closed-form price and sensitivity approximations for European swaptions allowed us to avoid cumbersome nested simulations that are associated with a naive bump-and-reval approach. Moreover, this enabled us to accurately estimate high dimensional, bucketed Delta- and Vega sensitivities, despite the fact that the risk-factors may be embedded in a low dimensional model. This yielded a richer handle for sensitivity modeling compared to, for example, a parallel shift in the model variables. Through several representative numerical examples, we demonstrated the performance of our approach under a one-factor model, benchmarked to the established least-square Monte Carlo method of Longstaff and Schwartz (2001). Overall, we observed superior convergence for the static replication method with regard to exposures, sensitivities and IM statistics. Thus, we demonstrated static replication to be a suitable alternative for exposure, dynamic IM and MVA quantification.

Disclosure statement

The opinions expressed in this work are solely those of the authors and do not represent in any way those of their current and past employers. No potential conflict of interest was reported by the authors.

Funding

This project has received funding from the NWO under the Industrial Doctorates grant.

References

- Andersen, L.B.G. and Piterbarg, V.V., *Interest Rate Modeling, Volume II: Term Structure Models*, 2010 (Atlantic Financial Press: London, UK).
- Andersson, K. and Oosterlee, C.W., A deep learning approach for computations of exposure profiles for high-dimensional bermudan options. *Appl. Math. Comput.*, 2021, **408**, 126332.
- Antonov, A., Issakov, S. and McClelland, A., Efficient SIMM-MVA calculations for callable exotics. Available at SSRN 3040061, 2017.
- Bank for International Settlements, Counterparty credit risk definitions and terminology, 2019.
- Burgard, C. and Kjaer, M., Funding strategies, funding costs. *Risk*, 2013, **26**, 82.
- Capriotti, L., Fast Greeks by algorithmic differentiation. *J. Comput. Finance*, 2011, **14**, 3.
- Capriotti, L., Jiang, Y. and Macrina, A., AAD and least-square Monte Carlo: Fast bermudan-style options and xVA Greeks. *Algorithmic Finance*, 2017, **6**, 35–49.
- Carriere, J.F., Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insur. Math. Econ.*, 1996, **19**, 19–30.
- Caspers, P. and Lichters, R., Initial margin forecast-bermudan swaption methodology and case study. Available at SSRN 3132008, 2018.
- De Graaf, C.S., Feng, Q., Kandhai, D. and Oosterlee, C.W., Efficient computation of exposure profiles for counterparty credit risk. *Int. J. Theor. Appl. Finance*, 2014, **17**, 1450024.
- Duffie, D. and Kan, R., A yield-factor model of interest rates. *Math. Financ.*, 1996, **6**, 379–406.
- Feng, Q., Jain, S., Karlsson, P., Kandhai, D. and Oosterlee, C.W., Efficient computation of exposure profiles on real-world and risk-neutral scenarios for bermudan swaptions. Available at SSRN 2790874, 2016.
- Filipovic, D., *Term-Structure Models. A Graduate Course*, 2009 (Springer: Berlin, Heidelberg).
- Fries, C.P., Fast stochastic forward sensitivities in Monte Carlo simulations using stochastic automatic differentiation (with applications to initial margin valuation adjustments). *J. Comput. Finance*, 2019, **22**, 103–125.
- Giles, M. and Glasserman, P., Smoking adjoints: Fast Monte Carlo Greeks. *Risk*, 2006, **19**, 88–92.
- Glasserman, P., *Monte Carlo Methods in Financial Engineering*, Vol. 53, 2004 (Springer: Berlin, Heidelberg).
- Glasserman, P. and Yu, B., Simulation for American options: Regression now or regression later? In *Monte Carlo and Quasi-Monte Carlo Methods 2002: Proceedings of a Conference held at the National University of Singapore, Republic of Singapore, 25–28 November, 2002*, pp. 213–226, 2004, (Springer: Berlin, Heidelberg).
- Glau, K., Pachon, R. and Pötz, C., Speed-up credit exposure calculations for pricing and risk management. *Quant. Finance*, 2021, **21**, 481–499.
- Green, A., *XVA: Credit, Funding and Capital Valuation Adjustments*, 2015 (John Wiley & Sons: Hoboken, NJ).
- Green, A.D. and Kenyon, C., Mva: Initial margin valuation adjustment by replication and regression. Available at SSRN 2432281, 2015.
- Gregory, J., *The XVA Challenge: Counterparty Risk, Funding, Collateral, Capital and Initial Margin*, 2020 (John Wiley & Sons: Chichester, UK).
- Griewank, A. and Walther, A., *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, 2008 (SIAM: Philadelphia, USA).
- Harrison, J.M. and Pliska, S.R., Martingales and stochastic integrals in the theory of continuous trading. *Processes and Their Applications; DIFdel; Stoch. Process. Their. Appl.*, 1981, **11**, 215–260.
- Hoencamp, J., Jain, S. and Kandhai, D., A semi-static replication method for bermudan swaptions under an affine multi-factor model. *Risks*, 2023, **11**, 168.
- International Swaps and Derivatives Association, Inc, Isda SIMM: From principles to model specification, 2016.
- International Swaps and Derivatives Association, Inc, Isda SIMM: Methodology, version 2.3, 2020.
- International Swaps and Derivatives Association, Inc, Isda margin survey year-end 2021, 2021.
- Jain, S. and Oosterlee, C.W., The stochastic grid bundling method: Efficient pricing of bermudan options and their Greeks. *Appl. Math. Comput.*, 2015, **269**, 412–431.
- Jain, S., Leitao, Á. and Oosterlee, C.W., Rolling adjoints: Fast Greeks along monte carlo scenarios for early-exercise options. *J. Comput. Sci.*, 2019, **33**, 95–112.
- Jamshidian, F., Libor and swap market models and measures. *Finance Stoch.*, 1997, **1**, 293–330.
- Jeanblanc, M. and Li, L., Characteristics and constructions of default times. *SIAM J. Financ. Math.*, 2020, **11**, 720–749.
- Joshi, M. and Kwon, O.K., Least squares monte carlo credit value adjustment with small and unidirectional bias. *Int. J. Theor. Appl. Finance*, 2016, **19**, 1650048.
- Kappen, C., Computing MVA via regression and principal component analysis, SSRN, 2017.
- Karlsson, P., Jain, S. and Oosterlee, C.W., Counterparty credit exposures for interest rate derivatives using the stochastic grid bundling method. *Appl. Math. Finance*, 2016, **23**, 175–196.
- Kingma, D.P. and Ba, J., Adam: A method for stochastic optimization, 2014, arXiv preprint arXiv:1412.6980.
- Lakhany, A. and Zhang, A., Efficient ISDA initial margin calculations using least squares monte-carlo, 2021, arXiv preprint arXiv:2110.13296.
- Lokeshwar, V., Bharadwaj, V. and Jain, S., Explainable neural network for pricing and universal static hedging of contingent claims. *Appl. Math. Comput.*, 2022, **417**, 126775.
- Longstaff, F.A. and Schwartz, E.S., Valuing American options by simulation: A simple least-squares approach. *Rev. Financ. Stud.*, 2001, **14**, 113–147.
- Lyashenko, A. and Mercurio, F., Looking forward to backward-looking rates: A modeling framework for term rates replacing LIBOR. Available at SSRN 3330240, 2019.
- Musiela, M. and Rutkowski, M., *Martingale Methods in Financial Modeling*, 2nd ed., 1997 (Springer: Berlin, Heidelberg).
- Schrager, D.F. and Pelsser, A.A., Pricing swaptions and coupon bond options in affine term structure models. *Math. Financ.*, 2006, **16**, 673–694.
- Shen, Y., Van Der Weide, J.A. and Anderluh, J.H., A benchmark approach of counterparty credit exposure of bermudan option under Lévy process: The Monte Carlo-cos method. *Procedia. Comput. Sci.*, 2013, **18**, 1163–1171.
- Simaitis, S., de Graaf, C., Hari, N. and Kandhai, D., Smile and default: The role of stochastic volatility and interest rates in counterparty credit risk. *Quant. Finance*, 2016, **16**, 1725–1740.
- Tsitsiklis, J.N. and Van Roy, B., Regression methods for pricing complex American-style options. *IEEE Trans. Neural Netw.*, 2001, **12**, 694–703.
- Zeron, M. and Ruiz, I., Dynamic initial margin via Chebyshev spectral decomposition, Tech. Rep., Working paper (24 August), 2018.
- Zhu, S.H. and Pykhtin, M., A guide to modeling counterparty credit risk, GARP Risk Review, July/August 2007.

Appendices

Appendix 1. Proof of Lemma 4.1

Proof Consider the process

$$y_t := \frac{P(t, T_1)}{P(t, T_0)}$$

As y represents a tradable asset denominated by the numéraire, it follows that y is a Martingale under \mathbb{Q}^{T_0} . An application of Itô's lemma and a subsequent measure change shows that y satisfies the following SDE

$$dy_t = y_t (B(t, T_0) - B(t, T_1))^\top \eta(t) dW_t^{T_0}$$

where W^{T_0} denotes a d -dimensional Brownian motion under \mathbb{Q}^{T_0} . Denoting $\hat{\eta}(t) := (B(t, T_0) - B(t, T_1))^\top \eta(t)$ and solving the SDE, conditioned on \mathcal{F}_t yields

$$P(T_0, T_1) := y_{T_0} = \frac{P(t, T_1)}{P(t, T_0)} e^{\int_t^{T_0} \hat{\eta}(u) dW_u^{T_0}}$$

Let $\tau_0 = T_0 - t$ and $\tau_1 = T_1 - t$. Note that by the definition of the zero-rate, the expression above can be rewritten as

$$P(T_0, T_1) = \exp \left\{ \tau_0 R(t, t + \tau_0) - \tau_1 R(t, t + \tau_1) + \int_t^{T_0} \hat{\eta}(u) dW_u^{T_0} \right\}$$

Therefore, the forward Delta sensitivity can be computed as

$$\frac{\partial P(T_0, T_1)}{\partial R_k(t)} = \left(\tau_0 \frac{\partial R(t, t + \tau_0)}{\partial R_k(t)} - \tau_1 \frac{\partial R(t, t + \tau_1)}{\partial R_k(t)} \right) P(T_0, T_1)$$

where $\alpha'_k(\tau_i) := \tau_i \frac{\partial R(t, t + \tau_i)}{\partial R_k(t)}$ denotes an \mathcal{F}_t -measurable coefficient. Under the assumption that the time- t zero-rate curve $\tau \mapsto R(t, t + \tau)$ is an interpolation between nodes $\{R_0(t), \dots, R_K(t)\}$, it follows that

$$\frac{\partial R(t, t + \tau_0)}{\partial R_k(t)} = \begin{cases} \frac{\tau - \tau_{k-1}}{\tau_{k+1} - \tau} & \text{if } \tau \in (\tau_{k-1}, \tau_k] \\ \frac{\tau_k - \tau_{k-1}}{\tau_{k+1} - \tau} & \text{if } \tau \in (\tau_k, \tau_{k+1}] \\ 0 & \text{otherwise} \end{cases}$$

This concludes the proof. \blacksquare

Appendix 2. Proof of Lemma 4.2

Proof Let $0 \leq t \leq T \leq T_m$. Note that we can write

$$\begin{aligned} \sigma_{\text{imp}}(T, T_m) &= \sqrt{\frac{1}{T_m - T} \int_T^{T_m} \|\tilde{\Sigma}_{m,M}^\top(u) \eta(u)\|_2^2 du} \\ &= \sqrt{\frac{1}{T_m - T} \left(\int_t^{T_m} \|\tilde{\Sigma}_{m,M}^\top(u) \eta(u)\|_2^2 du - \int_t^T \|\tilde{\Sigma}_{m,M}^\top(u) \eta(u)\|_2^2 du \right)} \\ &= \frac{1}{\sqrt{T_m - T}} \sqrt{(T_m - t) \tilde{\sigma}^2(t, T_m) - (T - t) \tilde{\sigma}^2(t, T)} \end{aligned}$$

Subsequently applying the chain rule to this expression yields

$$\begin{aligned} \frac{\partial \sigma_{\text{imp}}(T, T_m)}{\partial \sigma_k(t)} &= \frac{1}{\sqrt{T_m - T}} \frac{(T_m - t) \frac{\partial \tilde{\sigma}^2(t, T_m)}{\partial \sigma_k(t)} - (T - t) \frac{\partial \tilde{\sigma}^2(t, T)}{\partial \sigma_k(t)}}{2 \sqrt{(T_m - t) \tilde{\sigma}^2(t, T_m) - (T - t) \tilde{\sigma}^2(t, T)}} \\ &= \frac{(T_m - t) \tilde{\sigma}(t, T_m) \frac{\partial \tilde{\sigma}(t, T_m)}{\partial \sigma_k(t)} - (T - t) \tilde{\sigma}(t, T) \frac{\partial \tilde{\sigma}(t, T)}{\partial \sigma_k(t)}}{(T_m - T) \tilde{\sigma}_{\text{imp}}(T, T_m)} \end{aligned}$$

This concludes the proof. \blacksquare

Appendix 3. Proof of Lemma 4.3

Proof Recall that it is assumed that the instantaneous volatility in (1) takes values in \mathbb{R} and is piece-wise constant between dates $\{\tau_0, \dots, \tau_K\}$. Furthermore it is assumed that the volatility levels η_1, \dots, η_K are obtained through a bootstrap procedure, such that the implied volatility approximations $\tilde{\sigma}_{\text{imp}}$ given in (17) agree with the implied volatilities σ_k observed in the market. Lastly, note that only η_k and η_{k+1} are sensitive to a change in σ_k if other market-implied volatilities are kept fixed. For compactness of notation, we show the proof for $t = 0$. The generalization to $t > 0$ follows the same line of thought.

First consider the case $\tau_{k-1} \leq T < \tau_k$. Note that we can write

$$\begin{aligned} \tau_k \sigma_k^2 &= \int_0^{\tau_{k-1}} \Sigma_k^2(u) \eta^2(u) du + \eta_k^2 \int_{\tau_{k-1}}^{\tau_k} \Sigma_k^2(u) du \\ \implies \eta_k^2 &= \frac{\tau_k \sigma_k^2 - \int_0^{\tau_{k-1}} \Sigma_k^2(u) \eta^2(u) du}{\int_{\tau_{k-1}}^{\tau_k} \Sigma_k^2(u) du} \end{aligned} \quad (\text{A1})$$

Furthermore, we have the relation

$$\tilde{\sigma}(0, T) = \sqrt{\frac{1}{T} \left(\int_0^{\tau_{k-1}} \Sigma_{m,M}^2(u) \eta^2(u) du + \eta_k^2 \int_{\tau_{k-1}}^T \Sigma_{m,M}^2(u) du \right)}$$

It follows by the chain rule

$$\begin{aligned} \frac{\partial \tilde{\sigma}(0, T)}{\partial \sigma_k} &= \frac{1}{2 \tilde{\sigma}(0, T)} \frac{1}{T} \frac{\partial}{\partial \sigma_k} \\ &\quad \times \left(\int_0^{\tau_{k-1}} \Sigma_{m,M}^2(u) \eta^2(u) du + \eta_k^2 \int_{\tau_{k-1}}^T \Sigma_{m,M}^2(u) du \right) \\ &= \frac{1}{2T \tilde{\sigma}(0, T)} \left(\int_{\tau_{k-1}}^T \Sigma_{m,M}^2(u) du \right) \frac{\partial \eta_k^2}{\partial \sigma_k} \\ &= \frac{\tau_k \sigma_k}{T \tilde{\sigma}(0, T)} \frac{\int_{\tau_{k-1}}^T \Sigma_{m,M}^2(u) du}{\int_{\tau_{k-1}}^{\tau_k} \Sigma_k^2(u) du} \end{aligned}$$

where the last equality follows from substituting η_k^2 with (A1).

The case $\tau_k \leq T < \tau_{k+1}$ follows a similar line of thought. Note that we can write

$$\begin{aligned} \tau_{k+1} \sigma_{k+1}^2 &= \int_0^{\tau_{k-1}} \Sigma_{k+1}^2(u) \eta^2(u) du + \eta_k^2 \int_{\tau_{k-1}}^{\tau_k} \Sigma_{k+1}^2(u) du \\ &\quad + \eta_{k+1}^2 \int_{\tau_k}^{\tau_{k+1}} \Sigma_{k+1}^2(u) du \\ \tau_{k+1} \sigma_{k+1}^2 - \int_0^{\tau_{k-1}} \Sigma_{k+1}^2(u) \eta^2(u) du \\ \implies \eta_{k+1}^2 &= \frac{-\eta_k^2 \int_{\tau_{k-1}}^{\tau_k} \Sigma_{k+1}^2(u) du}{\int_{\tau_k}^{\tau_{k+1}} \Sigma_{k+1}^2(u) du} \end{aligned} \quad (\text{A2})$$

Furthermore, we have the relation

$$\tilde{\sigma}(0, T) = \sqrt{\frac{1}{T} \left(\int_0^{\tau_{k-1}} \Sigma_{m,M}^2(u) \eta^2(u) du + \eta_k^2 \int_{\tau_{k-1}}^{\tau_k} \Sigma_{m,M}^2(u) du + \eta_{k+1}^2 \int_{\tau_k}^T \Sigma_{m,M}^2(u) du \right)}$$

Applying the chain rule to the expression above, while substituting η_k^2 and η_{k+1}^2 with (A1) and (A2), respectively, we find

$$\begin{aligned} \frac{\partial \tilde{\sigma}(0, T)}{\partial \sigma_k} &= \frac{\tau_k \sigma_k}{T \tilde{\sigma}(0, T)} \\ &\quad \times \left(\frac{\int_{\tau_{k-1}}^{\tau_k} \Sigma_{m,M}^2(u) du}{\int_{\tau_{k-1}}^{\tau_k} \Sigma_k^2(u) du} - \frac{\int_{\tau_k}^T \Sigma_{m,M}^2(u) du}{\int_{\tau_k}^{\tau_{k+1}} \Sigma_{k+1}^2(u) du} \frac{\int_{\tau_{k-1}}^{\tau_k} \Sigma_{m,M}^2(u) du}{\int_{\tau_{k-1}}^{\tau_k} \Sigma_k^2(u) du} \right) \end{aligned}$$

If by approximation $\Sigma_{m,M} \approx \Sigma_k \approx \Sigma_{k+1}$, the expression above reduces to

$$\frac{\partial \tilde{\sigma}(0, T)}{\partial \sigma_k} = \frac{\tau_k \sigma_k}{T \tilde{\sigma}(0, T)} \frac{\int_T^{T_{k+1}} \Sigma_{m,M}^2(u) du}{\int_{T_k}^{T_{k+1}} \Sigma_{k+1}^2(u) du}$$

Appendix 4. Proof of Proposition 4.4

Proof For compactness, we will write $\boldsymbol{\pi}_m := \boldsymbol{\pi}_m(T_m)$. First, note that $\pi_{m,i}$ is continuously differentiable w.r.t. θ_t for all $i = 1, \dots, n$. It follows that λ_m is continuously differentiable w.r.t. both θ_t and \boldsymbol{w}_m . Secondly, we show that $D_w \lambda_m$ is non-singular. By the nature of the ReLU function, the payoffs $\pi_{m,1}, \dots, \pi_{m,n}$ are linearly independent when the strikes are disjoint. That is, there is no $\boldsymbol{v} \in \mathbb{R}^n$ such that $\boldsymbol{v}^\top \boldsymbol{\pi}_m = 0$ a.s. Now consider the matrix $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\pi}_m \boldsymbol{\pi}_m^\top]$. As a property of the auto-correlation matrix, it is known to be positive semi-definite, i.e.

$$\forall \boldsymbol{v} \in \mathbb{R}^n : \boldsymbol{v}^\top \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\pi}_m \boldsymbol{\pi}_m^\top] \boldsymbol{v} \geq 0$$

Through an argument by contradiction, we can show the matrix is in fact strictly positive definite. In that regard, assume there exists a $\boldsymbol{v} \in \mathbb{R}^n$ such that $\boldsymbol{v}^\top \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\pi}_m \boldsymbol{\pi}_m^\top] \boldsymbol{v} = 0$. It would follow that

$$0 = \boldsymbol{v}^\top \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\pi}_m \boldsymbol{\pi}_m^\top] \boldsymbol{v} = \mathbb{E}^{\mathbb{Q}}[\boldsymbol{v}^\top \boldsymbol{\pi}_m \boldsymbol{\pi}_m^\top \boldsymbol{v}] = \mathbb{E}^{\mathbb{Q}}[|\boldsymbol{v}^\top \boldsymbol{\pi}_m|^2]$$

■

this is in contradiction with the fact that $\pi_{m,1}, \dots, \pi_{m,n}$ are linearly independent. Hence we can conclude that $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\pi}_m \boldsymbol{\pi}_m^\top]$ is a real-valued, symmetric matrix, which is strictly positive definite. It follows that all its eigenvalues are real and positive. As the determinant of a symmetric matrix is the product of its eigenvalues, we can conclude that $D_w \lambda_m$ has a positive determinant and is non-singular. The result now follows by an application of the implicit function theorem. ■