

Accelerated Computations of Sensitivities for xVA

Griselda Deelstra^a, Lech A. Grzelak^{b,c}, Felix L. Wolf^{a,*}

^a*Department of Mathematics, Université libre de Bruxelles, Brussels, Belgium*

^b*Mathematical Institute, Utrecht University, Utrecht, the Netherlands*

^c*Rabobank, Utrecht, the Netherlands*

Abstract

Exposure simulations are fundamental to many xVA calculations and are a nested expectation problem where repeated portfolio valuations create a significant computational expense. Sensitivity calculations which require shocked and unshocked valuations in bump-and-revalue schemes exacerbate the computational load. A known reduction of the portfolio valuation cost is understood to be found in polynomial approximations, which we apply in this article to interest rate sensitivities of expected exposures. We consider a method based on the approximation of the shocked and unshocked valuation functions, as well as a novel approach in which the difference between these functions is approximated. Convergence results are shown, and we study the choice of interpolation nodes. Numerical experiments with interest rate derivatives are conducted to demonstrate the high accuracy and remarkable computational cost reduction. We further illustrate how the method can be extended to more general xVA models using the example of CVA with wrong-way risk.

Keywords: xVA, Sensitivity, Exposure Simulation, Portfolio Approximation, Polynomial Interpolation, Quadrature Nodes, Monte Carlo

1. Introduction

Value adjustments (xVA) are a core component in the risk management framework of modern asset pricing and are used to account for additional risks not covered under the standard theory of asset pricing. All xVA computations depend on the future evolution of the portfolios for which they are computed, so accurate exposure calculations are of high importance. In the presence of many stochastically modelled risk factors with potentially intricate dependencies, the exposure distribution quickly becomes intangible and warrants the use of Monte Carlo simulations. With a large number of simulation paths and evaluation times of the exposures, an enormous amount of required portfolio valuations may aggregate, which acts as a significant driver behind the often large computational demand of the exposure simulation.

One way to achieve faster simulations is through a reduction of the number of Monte Carlo paths or monitoring dates, but this is to the immediate detriment of accuracy. An alternative approach is to improve the speed of portfolio valuations.

Approximation techniques provide such an opportunity to reduce the computational complexity along the number of Monte Carlo simulations. At every time step, instead of evaluating the exact, and expensive, portfolio valuation function directly in

*Corresponding author

Email addresses: `Griselda.Deelstra@ulb.be` (Griselda Deelstra), `L.A.Grzelak@uu.nl` (Lech A. Grzelak), `Felix.Wolf@ulb.be` (Felix L. Wolf)

The views expressed in this paper are the personal views of the authors and do not necessarily reflect the views or policies of their current or past employers.

every realisation of its underlying risk factors, an accurate approximation is constructed from only a few exact evaluations. The resulting approximation can be evaluated very efficiently thanks to its polynomial structure. Similar approaches have been applied to improve efficiency in the simulation of differential equations (Xiu and Hesthaven, 2005) where they are known as stochastic collocation techniques. The approach considered in this article is closely connected to polynomial asset valuation approximation techniques, which have recently gained increased attention in the Mathematical Finance literature. Gaß et al. (2018) and Glau et al. (2019a) utilized polynomial interpolation over Chebyshev nodes for the pricing of options, respectively for the approximation of implied volatility functions. Glau et al. (2019b) introduced a novel approach to the pricing of American options, which uses Chebyshev interpolation to approximate the continuation values of the option. This approach allows for the splitting of the computation into a construction phase, where the approximation functions are computed, and a numerically efficient evaluation phase. In Glau et al. (2021), this dynamic approach is extended to a broader class of options, and to the nested expectation problem which is posed by exposure simulations. Here, it is also remarked that approximation techniques could prove suitable for sensitivity computations.

Grzelak et al. (2018) studied stochastic collocation techniques to improve the efficiency in sampling from random distributions, and showed optimality of quadrature interpolation nodes for (normally distributed) interest rates with respect to the L^2 -norm. In Grzelak (2022) the approximation of portfolio valuations with sparse grids in a multi-dimensional setting is considered.

Here, we extend the stochastic-collocation-based approach of exposure simulation to the calculation of expected exposure sensitivities at high numerical efficiency. We consider the calculation of interest-rate sensitivities, in particular the sensitivity with respect to market quotes used in the construction of the yield curve. These sensitivities are typically obtained from difference quotients, where the valuation under a particular market shock is compared to the valuation under normal, unshocked market conditions, so-called bump-and-revalue schemes (see, for example, Green (2015)).¹ For each market shock, such a scheme demands the full amount of exact portfolio valuations, which effectively multiplies the numerical effort by the number of market shocks considered.

Many attempts to manage this complexity are based on novel techniques, such as Adjoint Algorithmic Differentiation (Capriotti et al., 2017, Hüge and Savine, 2017) or neural-network-based xVA engines (Gnoatto et al., 2020). However, such methods require substantial overhead in their implementation, whereas the approach introduced in this article can be seamlessly integrated into standard Monte Carlo exposure simulations and merely requires the implementation of a polynomial interpolation scheme.

In this article, we introduce a new method for constructing estimations of the interest rate sensitivity of a portfolio’s expected exposure. We demonstrate how a polynomial interpolation technique can be directly applied to the unshocked and shocked portfolio valuation functions, leading to a significant reduction in computational requirements for sensitivity calculations. Moreover, we introduce a more efficient approach in which the difference between shocked and unshocked valuation functions is approximated with a reduced-degree polynomial, leading to increased complexity re-

¹In this article, “sensitivity” refers to the effect on the expected exposure induced by small changes to the term structure of interest rates. It can be thought of as a particular type of derivative, which is approximated by difference quotients. Then, “shocked” and “unshocked” valuations refer to the function valuations with and without shifted parameters, respectively. A comprehensive description of the procedure is provided in Section 3.

duction. We analyse convergence results and error bounds for the proposed methods and choices of interpolation nodes.

Our proposed approximation technique is compatible with sophisticated xVA frameworks, and we outline an adaptation to models that include wrong-way risk.

We conduct numerical experiments on portfolios of linear and non-linear assets, with a focus on large swap portfolios and Bermudan swaptions. In these experiments, we demonstrate how the L^2 -optimal quadrature nodes automatically resolve the challenge of finding a suitable interpolation domain.

The present article is structured as follows. In Section 2 we introduce exposure simulations in the classical sense and their approximations with polynomial interpolation. Section 3 summarizes the concept of interest-rate sensitivities, i.e., sensitivities with respect to the market quotes used in the construction of the yield curve, and defines the sensitivity of expected exposures. In Section 4, two approximation methods for these sensitivities are introduced. A *full-order approximation* applies the previously defined expected exposure approximations in a bump-and-revalue scheme. A *low-order difference approximation* directly approximates the functional change under the market shock. Section 5 contains an analytical error analysis, where we show convergence of the sensitivity approximation and also discuss the selection of interpolation nodes. In Section 6 we investigate the accuracy of the methods with numerical experiments, and in Section 7 we detail how the approach can be applied to Bermudan Swaptions with an additional numerical experiment. Finally, we summarize our findings in Section 8.

2. Exposure simulation and approximation

Before we introduce the sensitivity approximation methods studied in this article, we provide the context in which they arise, given by exposure simulations for xVA. We consider portfolios with underlying risk factors modelled by a stochastic process $(X(t))_{t \geq 0}$ on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$, where \mathbb{Q} is the risk-neutral measure. The stochastic process $X(t)$ takes values in \mathbb{R}^m , where m corresponds to the number of modelled risk factors. The portfolio value at time t is then described by a stochastic process $(V(t, X(t)))_{t \geq 0}$. When all future cash flows can be expressed in terms of payoffs $H(T_k, X(T_k))$, $k \in \{1, \dots, L\}$, $L \in \mathbb{N}$, the portfolio valuation is given by the conditional expectation²

$$V(t, X(t)) = \mathbb{E}_t^{\mathbb{Q}} \left[\sum_{k=1}^L \frac{B(t)}{B(T_k)} H(T_k, X(T_k)) \right], \quad (2.1)$$

where $B(t) = \exp(\int_0^t r(s) ds)$ is the numéraire of the measure \mathbb{Q} , i.e. we denote the risk-neutral interest (short) rate by $r(t)$.

In many xVA models, the valuation adjustment depends on whether a portfolio is associated with a cost or a benefit, that is, the models depend on the sign of the (future) valuations. This is expressed by the notion of positive and negative exposures, which are given by the positive and negative part of the valuation functions, $V^+(t, X(t))$ and $V^-(t, X(t))$, respectively. Exemplarily for credit valuation adjustments, a default of the counterparty only affects those portfolios which have a positive present value at the time of default. Equivalently from the point of view of the counterparty, its own default event is associated with its debit valuation adjustment, which analogously is based on the negative exposure (as seen by the counterparty).

²We denote expectations conditional on the filtration \mathcal{F}_t with a subscript \mathbb{E}_t .

The deterministic xVA quantities at initial time t_0 are based on *expected (positive or negative) exposures*, which are defined as the expectations of the respective discounted exposures. The expected positive³ exposure at monitoring date $t \geq t_0$ is given by

$$\text{EE}(t_0, t) = \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\frac{B(t_0)}{B(t)} V^+(t, X(t)) \right] \approx \frac{1}{M} \sum_{j=1}^M \frac{B(t_0)}{B(t; \omega_j)} V^+(t, X(t; \omega_j)), \quad (2.2)$$

which can be approximated by a Monte Carlo simulation approach with a large number of simulations $M \in \mathbb{N}$ of the underlying risk factors $X(t; \omega_j)$, $\omega_j \in \Omega$, $j \in \{1, \dots, M\}$.

Remark 2.1. *In this context, our focus is on the risk-neutral equivalent martingale measure \mathbb{Q} , which is suitable for simulating exposures for xVA pricing purposes. Nevertheless, the approach presented in this article is not limited to this measure and can also be used for computations under alternative measures such as the historical measure, or even technical choices such as the forward, swap, and other measures.*

2.1. Construction of approximating portfolio valuation functions

Given $R \in \mathbb{N}$ monitoring dates and M simulation paths, the Monte Carlo approximation in (2.2) requires $M \times R$ valuations of the portfolio. Each of these exact portfolio valuations may correspond to a contrived function evaluation, potentially due to the number of assets in the portfolio, or due to the presence of (non-linear) assets which require complex solution schemes. Therefore, the large number of exact portfolio valuations contributes an immense numerical expense to the exposure simulation. However, at each monitoring date t , the shape of the portfolio valuation function $V(t, \cdot)$ often exhibits a high degree of regularity when seen as a surface $\mathbb{R}^m \rightarrow \mathbb{R}$. This suggests that alternative functions $g(t, \cdot)$ exist, which generate almost indistinguishable valuation surfaces but can be numerically evaluated at much higher efficiency. Replacing the exact portfolio valuations with such approximation functions drastically reduces the computational burden of exposure simulations.

To obtain such approximating functions $g(t, \cdot)$ of the portfolio valuation functions $V(t, \cdot)$, we administer polynomial interpolation methods. The construction of a polynomial approximation based on N interpolation points $\{(x_k(t), V(t, x_k(t))) : k \in \{1, \dots, N\}\}$ requires only N of the expensive, exact portfolio valuations, where N is typically by orders of magnitude smaller than the number M of exact portfolio valuations utilized in the classical Monte Carlo approach (2.2). One of the challenges in the approach is to find appropriate interpolation nodes $x_k(t) \in \mathbb{R}^m$ in the domain of the underlying risk factor $X(t)$.

For the remainder of the article, we will focus on the base case of $m = 1$, i.e. portfolios with a single, one-dimensional risk factor. Particularly, in the numerical experiments of Section 6 and Section 7, we will consider interest rate products whose price at time t depends on the state of the underlying risk-free rate $r(t)$. The approximating functions $g(t, \cdot)$ can then be explicitly given in the Lagrange form

$$g(t, x) := \sum_{k=1}^N \left(V(t, x_k(t)) \prod_{\substack{\ell=1 \\ \ell \neq k}}^N \frac{x - x_\ell(t)}{x_k(t) - x_\ell(t)} \right) \approx V(t, x), \quad (2.3)$$

³The expected negative exposure is identical to (2.2) with the positive exposure replaced by the negative exposure $V^-(t, X(t))$. Within this article, we consider only the positive exposure, by symmetry the results can be extended to negative exposure computations.

after a suitable choice of interpolation nodes $(x_1(t), \dots, x_N(t))$ has been chosen, and the corresponding exact valuations $(V(t, x_1(t)), \dots, V(t, x_N(t)))$ have been computed. We define the set of N interpolation nodes by $\mathbf{x}_N(t)$. The described polynomial of degree $N - 1$ through a fixed choice of N interpolation points is unique, hence a numerical implementation may equivalently utilize a different polynomial basis, or a numerically more stable form, like the barycentric interpolation formula. We refer to Berrut and Trefethen (2004) for a treatise on its derivation and advantages.

Remark 2.2. Equation (2.3) sheds light on the required level of regularity for the portfolio valuation V . Specifically, the Lagrange remainder $R(t, x) := V(t, x) - g(t, x)$ of the polynomial interpolation g in N points depends on the N th derivative of the target function. It holds, for some value $\xi \in [\min_{1 \leq k \leq N} (x_k), \max_{1 \leq k \leq N} (x_k)]$, that

$$R(t, x) = \frac{\left. \frac{d^N V(t, x)}{dx^N} \right|_{x=\xi}}{N!} \prod_{k=1}^N (x - x_k(t)). \quad (2.4)$$

This suggests that the method may not be applicable if the N th derivative cannot be controlled. Conversely, a bounded N th derivative guarantee a certain quality of approximation over the domain spanned by the interpolation nodes. In Section 5, we extend this notion for expected exposures to the entire domain \mathbb{R} of the interest rate risk factor.

Monte Carlo exposure simulations can be greatly simplified with a stochastic collocation method, where exact valuations $V(t, \cdot)$ in (2.2) are replaced by approximations $g(t, \cdot)$. This yields an approximation of the expected (positive) exposure.

Definition 2.1 (Expected exposure approximation). *At each monitoring date t , let $V(t, \cdot): \mathbb{R} \rightarrow \mathbb{R}$ be a portfolio valuation function with stochastic risk factor $X(t)$. Let $g(t, \cdot)$ be an approximation of the portfolio valuation. Then, the expected (positive) exposure approximation is given by*

$$\widetilde{\mathbb{E}\mathbb{E}}(t_0, t) = \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\frac{B(t_0)}{B(t)} g^+(t, X(t)) \right] \approx \frac{1}{M} \sum_{j=1}^M \frac{B(t_0)}{B(t; \omega_j)} g^+(t, X(t; \omega_j)), \quad (2.5)$$

with $M \in \mathbb{N}$ samples of the risk factor.

Notably, at every monitoring date t , each approximated expected exposure $\widetilde{\mathbb{E}\mathbb{E}}(t_0, t)$ requires only N (depending on the degree of the polynomials $g(t, \cdot)$) exact, expensive portfolio valuations and M ‘cheap’ valuations of the polynomial approximation, which can be undertaken at high numerical efficiency. This stands in contrast to the M exact, expensive portfolio valuations in the classical approach (2.2).

Remark 2.3. *The approximation technique is not limited to basic xVA models that rely on isolated expected exposure terms. In general, the methods introduced here can be directly applied to many advanced xVA models, with the only requirement that explicit portfolio valuations take place within a Monte Carlo simulation. We outline the generalization on the example of unilateral CVA with wrong-way risk in Appendix A.*

In the remainder of the article, we adopt the expected-exposure form of Definition 2.1 which suffices to demonstrate the proposed methodology and has a straightforward representation.

3. Sensitivities with respect to the yield curve

We now introduce the concept of sensitivity with respect to the yield curve and apply it to the expected exposures previously defined.

In order to appropriately define the involved quantities and sensitivities, we briefly outline the yield curve construction. The yield curve represents the term structure of interest rates, which makes it a crucial tool in asset valuation.

Definition 3.1 (Yield curve and zero-coupon bond curve). *The yield curve Y of an investment horizon $[t_0, T]$ maps maturities to their associated total rate of return. That is, a payment of 1 (unit of the currency associated with the yield curve) at time t_0 returns an expected payment of $\exp(Y(S)(S - t_0))$ at time S . This is described by a real function*

$$Y: [t_0, T] \rightarrow \mathbb{R}, S \mapsto Y(S). \quad (3.1)$$

Under no-arbitrage conditions, the yield curve is directly related to the zero-coupon bond curve $P(t_0, \cdot)$, which indicates the equivalent discounted value that a payment of 1 at time S has at time t_0 ,

$$P(t_0, \cdot): [t_0, T] \rightarrow \mathbb{R}^+, S \mapsto P(t_0, S) := \exp(-Y(S)(S - t_0)). \quad (3.2)$$

A yield curve is calibrated to the market on the basis of a set of liquidly traded instruments with associated market quotes. We refer to this basis as the *market instruments*. Yield curve calibration refers to the process of finding a curve that reproduces each of these market quotes when the model prices of the instruments are calculated with respect to the curve. Typically, each instrument will only contribute information towards distinct parts of the yield curve. To obtain the entire curve over a continuous spectrum, suitable curve-fitting schemes must be applied, which affect the resulting yield curve. One such scheme is the multi-dimensional Newton–Raphson algorithm, which is applied in Oosterlee and Grzelak (2019) to obtain a discrete set of points on the yield curve. The effect that various interpolation methods between such points have on the yield curve is studied in Hagan and West (2006).

In this article, we use *sensitivity to the yield curve* as a general moniker, also the term *interest rate sensitivity* is known in the literature. To be precise, we refer to sensitivity with respect to the market quote of one of the market instruments from which the yield curve is constructed. Sensitivity refers to the rate of change, in our considerations of the expected exposure, to a change in the aforementioned market quote. In other words, it is a difference quotient which approximates the derivative of the expected exposure with respect to the market quote considered.

Each market instrument is associated with a particular maturity and particularly influences the yield curve in the surrounding region. Unlike approaches where only the finalized yield curve is shifted in a neighbourhood of the desired maturity, the sensitivity with respect to a market quote ensures that potential non-localities of the yield curve interpolation are correctly taken into account. Since the effect of different market quotes are of interest, the sensitivity with respect to the yield curve is rather a family of sensitivities or a *sensitivity profile*.

We formalise this in the following definitions.

Definition 3.2 (Market instruments). *The market instruments of a yield curve Y are a set of market instruments*

$$\Phi = \{\varphi_1, \dots, \varphi_n\}, \quad (3.3)$$

where at time t_0 , the market instrument φ_i , $i \in \{1, \dots, n\}$, corresponds to a maturity $T_i \in [t_0, T]$, has an available market quote K_i and a valuation function in terms of the curve Y .

In the following, we define the sensitivity of the expected exposure of a portfolio in terms of difference quotients obtained from *shocked* and *unshocked* valuations. Shocked valuations are the result of valuations using a shocked yield curve, which in turn is the yield curve obtained when a small shock, i.e. a linear shift, is applied to the quote of one of its market instruments.

Definition 3.3 (Shocked yield curves). *Let Y be a yield curve obtained from the market instruments Φ with market quotes $\{K_1, \dots, K_n\}$. For any $i \in \{1, \dots, n\}$, let Y_i be the yield curve obtained from the same set of instruments Φ , where the i -th instrument has its market quote shifted by a value $\Delta K_i \in \mathbb{R}$. That is, the market quotes of yield curve Y_i 's market instruments are $\{K_1, \dots, K_{i-1}, K_i + \Delta K_i, K_{i+1}, \dots, K_n\}$.*

We can now define the precise sensitivities considered in this article.

Definition 3.4 (Sensitivity of the expected exposure with respect to the yield curve). *Let $V(t; Y)$ denote the portfolio valuation at time $t \in [t_0, T]$ under the assumption of an initial yield curve Y . For any $i \in \{1, \dots, n\}$, the sensitivity of V with respect to K_i is defined by*

$$\frac{\partial V(t; Y)}{\partial K_i} := \lim_{\Delta K_i \rightarrow 0} \frac{V(t; Y_i) - V(t; Y)}{\Delta K_i}. \quad (3.4)$$

Let $\text{EE}(t_0, t)$ be the expected (positive) exposure of V as given in (2.2). The sensitivity of the expected exposure at time t with respect to the market quote K_i is given by

$$\frac{\partial \text{EE}(t_0, t)}{\partial K_i} = \frac{\partial}{\partial K_i} \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\frac{B(t_0)}{B(t)} V^+(t, X(t)) \right]. \quad (3.5)$$

The sensitivity of the expected exposure can be rewritten by exchanging the order of the expectation and differentiation operator, followed by the product rule of differentiation. This yields the formula

$$\begin{aligned} \frac{\partial \text{EE}(t_0, t)}{\partial K_i} &= \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\frac{\partial}{\partial K_i} \left(\frac{B(t_0)}{B(t)} V^+(t, X(t)) \right) \right] \\ &= \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) V^+(t, X(t)) + \frac{B(t_0)}{B(t)} \frac{\partial V^+(t, X(t))}{\partial K_i} \right]. \end{aligned} \quad (3.6)$$

This expanded form is suitable to approximation in the common framework of Monte Carlo exposure simulations and allows us to relate the sensitivity of the expected exposure to sensitivities of the numéraire and of the portfolio.

The derivatives in (3.6) can not be assumed to be analytically available, depending on the construction methodology of the yield curve and the type of assets in the portfolio. Therefore, we will consider forward differences for some fixed, small values of $\Delta K_i > 0$ in (3.4).

We introduce some short-hand notations. In the unshocked market, we continue denoting the interest rate by $r(t)$, the money market account by $B(t)$, the risk factor by $X(t)$ and the value of the portfolio by $V(t, X(t))$. In the i -th shocked market, we denote the interest rate by $r_i(t)$, the money market account by $B_i(t) = \exp(\int_0^t r_i(s) ds)$, the risk factor by $X_i(t)$ and the portfolio valuation by $V_i(t, X_i(t))$. We emphasize that the market shock affects the entire discounting structure.

The resulting Monte Carlo simulation scheme for the sensitivity of an interest rate asset with M paths ω_j , $j = 1, \dots, M$, is given by

$$\begin{aligned} \frac{\partial}{\partial K_i} \mathbb{E}\mathbb{E}(t_0, t) &\approx \sum_{j=1}^M \left(\frac{\partial}{\partial K_i} \frac{B(t_0; \omega_j)}{B(t; \omega_j)} \right) V^+(t, X(t; \omega_j)) + \frac{B(t_0; \omega_j)}{B(t; \omega_j)} \frac{\partial V^+(t, X(t; \omega_j))}{\partial K_i} \\ &\approx \sum_{j=1}^M \frac{\frac{B_i(t_0; \omega_j)}{B_i(t; \omega_j)} - \frac{B(t_0; \omega_j)}{B(t; \omega_j)}}{\Delta K_i} V^+(t, X(t; \omega_j)) + \frac{B(t_0; \omega_j)}{B(t; \omega_j)} \frac{V_i^+(t, X_i(t; \omega_j)) - V^+(t, X(t; \omega_j))}{\Delta K_i}. \end{aligned} \quad (3.7)$$

This is equivalent to the forward sensitivity approach given in Green (2015).

4. Approximation of expected exposure sensitivities

Analogous to the expected exposure approximation in Definition 2.1, we define two approximations of the expected exposure sensitivity (3.6). First, a *full-order approximation*, in which both the shocked and unshocked portfolio valuation functions, $V(t, \cdot)$ and $V_i(t, \cdot)$ are approximated with polynomials of full degree $N - 1$, i.e. constructed with N exact portfolio valuations each. Secondly, we define a *low-order difference approximation*, where only $V(t, \cdot)$ is approximated at full degree, and the approximations of $V_i(t, \cdot)$ are constructed with a lower degree approximation of the difference between the shocked and unshocked portfolio valuation functions.

Definition 4.1 (The full-order approximation of the sensitivity). *For any monitoring date $t_0 \leq t \leq T$, let the exact portfolio valuation $V(t, \cdot)$ associated with the yield curve Y be approximated by the polynomial $g(t, \cdot)$, based on N interpolation nodes $\mathbf{x}_N(t)$, given in (2.3).*

For a fixed $i \in \{1, \dots, n\}$, let $V_i(t, \cdot)$ be the exact portfolio valuation function associated with the yield curve Y_i of the i -th market shock, given in Definition 3.3. Let $g_i(t, \cdot)$ be a polynomial approximation function of $V_i(t, \cdot)$ based on N interpolation nodes $(x_1^i(t), \dots, x_N^i(t))$, the set of which we denote by $\mathbf{x}_N^i(t)$.

The full-order approximation of the sensitivity of the expected exposure with respect to market quote K_i , $\Psi_N^i(t) \approx \partial \mathbb{E}\mathbb{E}(t_0, t) / \partial K_i$, is given by

$$\begin{aligned} \Psi_N^i(t) &:= \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) g^+(t, X(t)) + \frac{B(t_0)}{B(t)} \frac{g_i^+(t, X_i(t)) - g^+(t, X(t))}{\Delta K_i} \right] \\ &\approx \sum_{j=1}^M \frac{\frac{B_i(t_0; \omega_j)}{B_i(t; \omega_j)} - \frac{B(t_0; \omega_j)}{B(t; \omega_j)}}{\Delta K_i} g^+(t, X(t; \omega_j)) + \frac{B(t_0; \omega_j)}{B(t; \omega_j)} \frac{g_i^+(t, X_i(t; \omega_j)) - g^+(t, X(t; \omega_j))}{\Delta K_i}. \end{aligned}$$

Remark 4.1. *There are two sources of approximation error in the above definition. First, the accuracy of the approximation g hinges on the regularity of the portfolio valuation function V , as detailed in Remark 2.2. This approximation error may potentially be addressed by increasing the number of interpolation nodes N .*

Secondly, the construction of the yield curve itself may introduce inaccuracies. In Appendix C, we show connections between the sensitivity of the yield curve and its modelling methodology, which in turn may affect the quality of the approximation. This source of inaccuracy requires refinement of the yield curve construction methodology and is not directly related to the interpolation of the portfolio valuation function.

The full-order approximation greatly reduces the number of exact portfolio valuations. In the classical approach given in (3.7), $2M$ exact portfolio valuations are required, where we recall that M is the number of paths in the Monte Carlo simulation. In the full-order approach, this is reduced to $2N$ exact valuations to obtain the approximation functions, and $2M$ (cheap) valuations of the polynomial approximations. The full-order approach to sensitivity in Definition 4.1 relies on polynomial approximations based on N exact valuations at their interpolation nodes. In the computation of a sensitivity profile with respect to multiple market quotes $\{K_1, \dots, K_n\}$, each sensitivity approximation reuses the same approximation $g(t, \cdot)$ of $V(t, \cdot)$ but requires a newly constructed approximation $g_i(t, \cdot)$ of $V_i(t, \cdot)$. Thus, n different market shocks require $(n + 1)N$ exact portfolio valuations per monitoring date.

We now introduce a low-order difference approximation which reduces the number of exact valuations needed to approximate the shocked portfolio valuation $V_i(t, \cdot)$.

In the low-order difference approach, for every monitoring date t , an approximation $\tilde{g}_i(t, \cdot) \approx V_i(t, \cdot)$ is constructed based on the previously established estimator of the unshocked portfolio function $g(t, \cdot)$, to which another polynomial $h_i(t, \cdot)$ of lower degree is added. That is, we set

$$\tilde{g}_i(t, \cdot) := g(t, \cdot) + h_i(t, \cdot). \quad (4.1)$$

By constructing $h_i(t, \cdot)$ with $d < N$ interpolation nodes, only d exact portfolio valuations $V^i(t, \cdot)$ are required, and consequently, the total number of exact portfolio valuations in the sensitivity profile is reduced to $nd + N$ at each monitoring date.

We construct $h_i(t, \cdot)$ as a polynomial of degree $d - 1$ which approximates the difference between the polynomial approximation $g(t, \cdot)$ and the shocked portfolio, $V_i(t, \cdot)$. That is, in Lagrange form, the difference approximation $h_i(t, \cdot)$ is given by

$$h_i(t, x) := \sum_{j=1}^d (V_i(t, x_{j,d}^i(t)) - g(x_{j,d}^i(t))) L_{j,d,t}^i(x), \quad (4.2)$$

with Lagrange basis

$$L_{j,d,t}^i(x) := \prod_{\substack{k=1 \\ k \neq j}}^d \frac{x - x_{k,d}^i(t)}{x_{j,d}^i(t) - x_{k,d}^i(t)}, \quad (4.3)$$

where $(x_{1,d}^i(t), \dots, x_{d,d}^i(t))$ are the interpolation nodes used in the construction of $h_i(t, \cdot)$ which we collect in the set $\mathbf{x}_d^i(t)$.

We now define the low-order difference approximation of the sensitivity.

Definition 4.2 (The low-order difference approximation of the sensitivity). *For any monitoring date $t_0 \leq t \leq T$ and shock to market quote K_i , $i \in \{1, \dots, n\}$, let the unshocked portfolio valuation function $V(t, \cdot)$, its approximation $g(t, \cdot)$, and the shocked portfolio valuation function $V_i(t, \cdot)$ be defined as in Definition 4.1. Let $\tilde{g}_i(t, \cdot)$ be the approximation of $V_i(t, \cdot)$ given in (4.1) and (4.2). The low-order difference approximation of the sensitivity of the expected exposure with respect to market quote K_i , $\Psi_{d,N}^i(t) \approx \partial \text{EE}(t_0, t) / \partial K_i$, is given by*

$$\Psi_{d,N}^i(t) := \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) g^+(t, X(t)) + \frac{B(t_0)}{B(t)} \frac{\tilde{g}_i^+(t, X_i(t)) - g^+(t, X(t))}{\Delta K_i} \right]. \quad (4.4)$$

When the set of interpolation nodes $\mathbf{x}_d^i(t)$ of $h_i(t, \cdot)$ are chosen as a subset of the N interpolation nodes $\mathbf{x}_N^i(t)$ on which the approximation $g_i(t, \cdot)$ in the full-order approach

is based, another interpretation of the low-order difference method arises. Under this particular choice of interpolation nodes, it becomes possible to rewrite (4.2) as

$$\begin{aligned} h_i(t, x) &= \sum_{j=1}^d V_i(t, x_{j,d}^i(t)) L_{j,d,t}^i(x) - \sum_{j=1}^d g(t, x_{j,d}^i(t)) L_{j,d,t}^i(x) \\ &= \sum_{j=1}^d g_i(t, x_{j,d}^i(t)) L_{j,d,t}^i(x) - \sum_{j=1}^d g(t, x_{j,d}^i(t)) L_{j,d,t}^i(x), \end{aligned} \quad (4.5)$$

since g_i and V_i coincide in the interpolation nodes $x_{j,d}^i(t)$. Thus, we may relate $h_i(t)$ solely to the approximations g_i and g . In particular, we may interpret h_i^t as the difference between a polynomial approximation $p_i(t, \cdot)$ of $g_i(t, \cdot)$, and a polynomial approximation $p(t, \cdot)$ of $g(t, \cdot)$, both of reduced degree $d - 1$. Summarizing,

$$h_i(t, x) =: p_i(t, x) - p(t, x). \quad (4.6)$$

With this representation at hand, it becomes easy to see that if the difference approximation is constructed with the full number of $d = N$ points, then, by the uniqueness of the polynomial approximation, it holds that $p_i(t, \cdot) = g_i(t, \cdot)$ and $p(t, \cdot) = g(t, \cdot)$. Consequently, the approximations $\tilde{g}_i(t, \cdot)$ and $g_i(t, \cdot)$ must coincide and the same holds for the expected exposure sensitivity estimators $\Psi_{d,N}^i(t)$ and $\Psi_N^i(t)$, when $d = N$.

5. Error Analysis

We have thus introduced two approximators of the expected exposure's yield curve sensitivity, the full-order estimator $\Psi_N^i(t)$ given in Definition 4.1 and the low-order difference estimator $\Psi_{d,N}^i(t)$ given in Definition 4.2. The benchmark of the error analysis is provided by the conventional estimator, the finite difference approach with exact portfolio valuations given by

$$\Psi^i(t) := \mathbb{E}_{t_0} \left[\left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) V^+(t, X(t)) + \frac{B(t_0)}{B(t)} \frac{V_i^+(t, X(t)) - V^+(t, X(t))}{\Delta K_i} \right]. \quad (5.1)$$

For every time t , we analyse the absolute approximation error between the conventional estimator and the low-order difference approximation. Since it holds that $\Psi_{d,N}^i(t) = \Psi_N^i(t)$ for the special case of $d = N$, the resulting error bound can be quickly extended to the full-order approach. The absolute error is given by

$$\begin{aligned} |\Psi^i(t) - \Psi_{d,N}^i(t)| &= \left| \mathbb{E}_{t_0} \left[\left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) (V^+(t, X(t)) - g^+(t, X(t))) \right. \right. \\ &\quad \left. \left. + \frac{B(t_0)}{B(t)} \left(\frac{V_i^+(t, X_i(t)) - V^+(t, X(t))}{\Delta K_i} - \frac{\tilde{g}_i^+(t, X_i(t)) - g^+(t, X(t))}{\Delta K_i} \right) \right] \right|. \end{aligned}$$

In the following, we will not write the second argument of the functions V, V_i, g, g_i when using $L^q(\Omega)$ norms for $q \in \{1, 2\}$. Observe that for any two arbitrary functions f_1, f_2 , it holds that

$$|f_1^+(x) - f_2^+(x)| \leq |f_1(x) - f_2(x)|. \quad (5.2)$$

Together with an application of Jensen's inequality and the triangle inequality, we obtain a first upper bound

$$\begin{aligned} |\Psi^i(t) - \Psi_{d,N}^i(t)| &\leq \mathbb{E}_{t_0} \left[\left\| \left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) (V(t, X(t)) - g(t, X(t))) \right\| \right] \\ &\quad + \mathbb{E}_{t_0} \left[\left\| \frac{B(t_0)}{B(t)} \frac{V_i(t, X_i(t)) - \tilde{g}_i(t, X_i(t))}{\Delta K_i} \right\| \right] + \mathbb{E}_{t_0} \left[\left\| \frac{B(t_0)}{B(t)} \frac{V(t, X(t)) - g(t, X(t))}{\Delta K_i} \right\| \right] \\ &= \left\| \left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) (V(t) - g(t)) \right\|_1 + \left\| \frac{B(t_0)}{B(t)} \frac{V_i(t) - \tilde{g}_i(t)}{\Delta K_i} \right\|_1 + \left\| \frac{B(t_0)}{B(t)} \frac{V(t) - g(t)}{\Delta K_i} \right\|_1. \end{aligned}$$

An application of Hölder's inequality allows us to relate the absolute error directly to distances between target functions and their approximations,

$$\begin{aligned} |\Psi^i(t) - \Psi_{d,N}^i(t)| &\leq \left\| \left(\frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right) \right\|_2 \|V(t) - g(t)\|_2 \\ &\quad + \left| \frac{1}{\Delta K_i} \right| \left\| \frac{B(t_0)}{B(t)} \right\|_2 (\|V_i(t) - \tilde{g}_i(t)\|_2 + \|V(t) - g(t)\|_2). \end{aligned}$$

The terms related to the discount factor in this inequality are quickly bounded. Because of its connection to the zero-coupon bond price $P(t_0, t)$, we immediately obtain that the L^2 -norm of the discount factor is finite, and denote it by

$$C_1(t) := \left\| \frac{B(t_0)}{B(t)} \right\|_2 = \mathbb{E}_{t_0} \left[\exp \left(-2 \int_{t_0}^t r(s) ds \right) \right]^{\frac{1}{2}} < \infty. \quad (5.3)$$

A bound on the norm of the derivative of the discount factor is encoded in the assumptions of the yield curve construction. Generally, it holds that

$$\begin{aligned} C_2(t) &:= \left\| \frac{\partial}{\partial K_i} \frac{B(t_0)}{B(t)} \right\|_2 = \mathbb{E}_{t_0} \left[\left(\frac{\partial}{\partial K_i} \exp \left(- \int_{t_0}^t r(s) ds \right) \right)^2 \right]^{\frac{1}{2}} \\ &= \mathbb{E}_{t_0} \left[\left(\int_{t_0}^t \frac{\partial}{\partial K_i} r(s) ds \right)^2 \exp \left(-2 \int_{t_0}^t r(s) ds \right) \right]^{\frac{1}{2}}. \quad (5.4) \end{aligned}$$

One of the desirable features of yield curve interpolation described by Hagan and West (2006) is *(forward) stability*, which describes how much the interpolated yield changes with respect to the inputs. The authors empirically confirm boundedness for a broad range of interpolation methods. Connecting forward stability to the interest rate becomes possible under concrete assumptions about the interest rate model used. In Appendix B, this is shown explicitly for the 1-factor Hull–White interest rate model. Thus, for suitable interest rate and yield curve models, we may assume a bound

$$|\Psi^i(t) - \Psi_{d,N}^i(t)| \leq C_2(t) \|V(t) - g(t)\|_2 + \frac{C_1(t)}{|\Delta K_i|} (\|V_i(t) - \tilde{g}_i(t)\|_2 + \|V(t) - g(t)\|_2),$$

where $C_1(t), C_2(t) \in \mathbb{R}$ are the constants in (5.3) and (5.4), independent of the number of nodes N and d .

In the following, we identify the case $i = 0$ with the market without parameter shocks, so that $V_0(t, x) = V(t, x)$ and $g_0(t, x) = g(t, x)$. Then, the approximation errors of the portfolio valuation function, respectively its shocked market versions, are given by

$$\varepsilon_i(t) := \|V_i(t) - g_i(t)\|_2, \quad i \in \{0, \dots, N\}. \quad (5.5)$$

Recall that $\mathbf{x}_N^i(t)$ are the N interpolation nodes used in the construction of $g_i(t)$ in the full-order approach and $\mathbf{x}_d^i(t)$ are the d interpolation nodes used in the construction of $\tilde{g}_i(t)$ in the low-order difference approach. In the previous section, it was shown that choosing the interpolation nodes of the low-order difference approach as a subset of the nodes in the full-order approach,

$$\mathbf{x}_d^i(t) \subseteq \mathbf{x}_N^i(t) \quad \text{for all } d \leq N, \quad (5.6)$$

allows for a convenient interpretation of the low-order difference polynomial $h_i(x)$. It was shown in (4.6), that the low-order difference polynomial can be expressed by $h_i(t, \cdot) = p_i(t, \cdot) - p(t, \cdot)$, a difference between the degree $d - 1$ approximations of $g_i(t, \cdot)$ and $g(t, \cdot)$. This allows for further bounds

$$\begin{aligned} \|V_i(t) - \tilde{g}_i(t)\|_2 &\leq \|V_i(t) - g_i(t)\|_2 + \|g_i(t) - \tilde{g}_i(t)\|_2 \\ &= \varepsilon_i(t) + \|g_i(t) - g(t) - (p_i(t) - p(t))\|_2 \\ &\leq \varepsilon_i(t) + \|g_i(t) - p_i(t)\|_2 + \|g(t) - p(t)\|_2. \end{aligned} \quad (5.7)$$

Analogously to (5.5), we set $p_0(t) := p(t)$ and denote the errors of approximating the degree $N - 1$ polynomials with their degree $d - 1$ counterparts by

$$\delta_i(t) := \|g_i(t) - p_i(t)\|_2, \quad i \in \{0, \dots, N\}, \quad (5.8)$$

so that, combining these results, we obtain the error bound of the low-degree approximation,

$$|\Psi^i(t) - \Psi_{d,N}^i(t)| \leq C_2(t)\varepsilon_0(t) + \frac{C_1(t)}{|\Delta K_i|}(\varepsilon_0(t) + \varepsilon_i(t) + \delta_0(t) + \delta_i(t)). \quad (5.9)$$

For $d = N$, it holds that $\delta_i(t) = 0$ and (5.9) provides the error bound for the full-order approximation. Its convergence follows in the next section by showing that $\varepsilon_i \rightarrow 0$ as $N \rightarrow \infty$.

5.1. Quadrature nodes

For $i \in \{0, \dots, n\}$, let $I_t^i := [\min(\mathbf{x}_N^i(t)), \max(\mathbf{x}_N^i(t))]$ denote the interval spanned by the interpolation nodes $\mathbf{x}_N^i(t) = \{x_1^i(t), \dots, x_N^i(t)\}$, where we set $\mathbf{x}_N^0(t) := \mathbf{x}_N(t)$. Furthermore, let $|I_t^i|$ denote its length. When the target function $V_i(t, x)$ is sufficiently regular, the accuracy of its polynomial interpolation $g_i(t, x)$ can be explicitly given, for every $x \in I_t^i$, in terms of the Lagrange bound,

$$V_i(t, x) - g_i(t, x) = \frac{\left. \frac{d^N V_i(t, x)}{dx^N} \right|_{x=\xi_i}}{N!} \prod_{k=1}^N (x - x_k^i(t)), \quad i \in \{0, \dots, n\}, \quad (5.10)$$

where the derivative of $V_i(t, x)$ is evaluated in some value $\xi_i \in I_t^i$. Since the approximation setting considered in this article does not permit control over the function $V_i(t, x)$ or its derivatives, we may only exert influence over the error of the interpolation through the choice of interpolation nodes $\mathbf{x}_N^i(t)$. The optimal choice of interpolation nodes depends on the norm under which $V_i(t, x) - g_i(t, x)$ is considered.

The $L^2(\Omega)$ norm in (5.5) implicitly depends on the random variable $X_i(t)$, where we again set $X_0(t) := X(t)$. When the risk-factor $X_i(t)$ admits a density $f_{X_i(t)}$, the error with respect to the $L^2(\Omega)$ norm can be expressed as an integral with weight function $f_{X_i(t)}(x)$, and Grzelak et al. (2018) showed that with respect to the $L^2(\Omega)$

norm required in (5.5), the optimal choice of interpolation nodes are the quadrature points $y_k^i(t)$ of the associated quadrature rule,

$$\begin{aligned}\varepsilon_i(t)^2 &= \int_{\mathbb{R}} (V_i(t, x) - g_i(t, x))^2 f_{X_i(t)}(x) dx \\ &= \sum_{k=1}^N (V_i(t, y_k^i(t)) - g_i(t, y_k^i(t)))^2 w_k^i(t) + \tilde{\varepsilon}_N^i(t).\end{aligned}\quad (5.11)$$

Here, $w_k^i(t)$ are the corresponding quadrature weights and $\tilde{\varepsilon}_N^i(t)$ is the quadrature error. The quadrature points $y_k^i(t)$ can be constructed with the recurrence relation of Golub and Welsch (1969), whenever the corresponding moments of $X_i(t)$ are available. Alternatively, when the orthogonal polynomial family of $L^2(\Omega)$ is known, they can be found as the zeros of these polynomials. Choosing the interpolation points in this manner as $x_k^i(t) = y_k^i(t)$ yields

$$\varepsilon_i(t)^2 = \tilde{\varepsilon}_N^i(t), \quad (5.12)$$

and convergence $\varepsilon_i(t) \rightarrow 0$, as $N \rightarrow \infty$, can be inferred from convergence of the respective quadrature rule, which depends both on the risk factor $X_i(t)$ and the valuation function $V_i(t)$.

In Section 6 and Section 7, we will consider interest rate products with normally distributed risk factors $X(t) = r(t)$. Then, the optimal interpolation points are given by the zeros of the Hermite polynomials, and the convergence of the quadrature error, as $N \rightarrow \infty$, is obtained for sufficiently differentiable portfolio valuation functions $V_i(t)$ from the convergence of the Gauss-Hermite quadrature error (Abramowitz and Stegun, 1964).

Of the products considered, interest rate swap yield a valuation function which is a linear combination of analytic functions and is thus easily seen as sufficiently smooth. Gaß et al. (2018) prove that the valuation function of European options is analytic, which can be applied to (European) swaptions. Glau et al. (2021) note that the valuation function of a Bermudan swaption is merely continuously differentiable, but also offer a practical method to determine convergence based on the decay of polynomial coefficients as the degree of the approximation function increases, a comparable approach is discussed in Trefethen (2019).

The approximation errors of the low-order contribution, $\delta_i(t)$, $i \in \{0, \dots, N\}$, given in (5.8), can be bounded over their domain of interpolation with the Lagrange error (5.10), where now the target functions $g_i(t)$ are polynomials and therefore analytic functions by construction. To ensure the nestedness of interpolation nodes $\mathbf{x}_d^i(t)$ for $d < N$, we choose them as the inner sets of $\mathbf{x}_N^i(t)$, where nodes are alternately removed from either end of the interval until the desired number of nodes remains,

$$\mathbf{x}_d^i(t) := \{x_{1+\lfloor \frac{N-d}{2} \rfloor}^i(t), \dots, x_{N-\lceil \frac{N-d}{2} \rceil}^i(t)\}. \quad (5.13)$$

5.2. Chebyshev nodes

An alternative selection of interpolation nodes is given by the Chebyshev nodes. These are given by $c_k := \cos(\frac{k\pi}{N})$ for the interval $[-1, 1]$, and can be scaled to the domain I_t^i with a linear transformation $T_t^i: [-1, 1] \rightarrow I_t^i$. It is well-known, see for example Burden et al. (2015), that Chebyshev nodes allow for a stricter bound of the Lagrange error over their interpolation domain. From this, it can be deduced that

with respect to the supremum norm it holds that

$$\|V_i(t, x) - g_i(t, x)\|_\infty \leq \frac{|I_t^i|^N}{2^N N!} \left\| \left. \frac{d^N V_i(t, x)}{dx^N} \right|_{x=\xi} \right\|_\infty. \quad (5.14)$$

In this norm, the choice of Chebyshev interpolation nodes $x_k^i(t) = T_t^i(c_k)$ is considered optimal and convergence $\|V_i(t, x) - g_i(t, x)\|_\infty \rightarrow 0$ for $N \rightarrow \infty$ is easily shown whenever $V_i(t, x)$ is sufficiently regular.

The supremum norm is easily interpreted as the maximal deviation of the approximation function $g_i(t, x)$ from its target $V_i(t, x)$, and it implies convergence with respect to many other norms. However, these results are limited to the necessarily finite domain of interpolation I_t^i .

6. Numerical experiments: interest-rate swaps

In this section, we conduct numerical experiments on portfolios of interest rate swaps, i.e., the underlying risk factor of the portfolio is $X(t) = r(t)$. In the first experiment, we explore the behaviour of the proposed methods with a small portfolio comprised of a single swap. Then, we move on to a larger portfolio, which consists of many different payer and receiver swaps.

In all experiments, we construct a yield curve based on the market instruments given in Appendix D, Table D.3. Each of these instruments is associated with a market quote K_i and a maturity T_i , $i \in \{1, \dots, 8\}$. Sensitivities with respect to the yield curve are based on individual shocks to each of these rates K_i , so that in total, $n = 8$ additional, shocked yield curves are constructed as delineated in Definition 3.3. As we will see shortly, there exists a correspondence between the maturity T_i of the shocked instrument, and the region of effect on the yield curve. This motivates a practise of referring to the shock of instrument K_i as the shock at maturity T_i .

To obtain a good approximation of the derivative (3.6), we fix the shock size to $\Delta K_i = 0.0001 = 1$ bp.

The interest rates $r(t)$, and its shocked-market versions $r_i(t)$, are modelled with 1-factor Hull–White models (Hull and White, 1990). In the following, we continue to identify the case $i = 0$ with the market without parameter shocks, so that $r_0(t) = r(t)$. The prescribed dynamics are

$$dr_i(t) = \lambda(\theta_i(t) - r(t))dt + \eta W(t), \quad 0 \leq i \leq n, \quad (6.1)$$

with a Brownian motion $W(t)$, speed of mean reversion $\lambda = 0.01$, volatility coefficient $\eta = 0.02$, and time-dependent drift term $\theta_i(t)$ chosen such that the bond prices obtained from the model fit the respective yield curves. Later on, in Section 6.1, we also consider a stressed scenario with an increased volatility coefficient.

At each time t , the interest rates follow normal distributions $r_i(t) \sim \mathcal{N}(\mu_i(t), \sigma_i^2(t))$ with explicitly known distribution parameters. As discussed in Section 5.1, we choose the interpolation nodes $\mathbf{x}_N^i(t)$ based on these distributions as the zeros of the (probabilist's) Hermite polynomials orthogonal under $L^2(\Omega)$, and the reduced interpolation nodes $\mathbf{x}_d^i(t)$ as their inner subsets given in (5.13).

6.1. Single swap portfolio

We consider a portfolio consisting of a single interest-rate payer swap with twice-yearly payments at times $\bar{T}_1, \dots, \bar{T}_m$, final maturity $\bar{T}_m = 20$ years, and a notional of

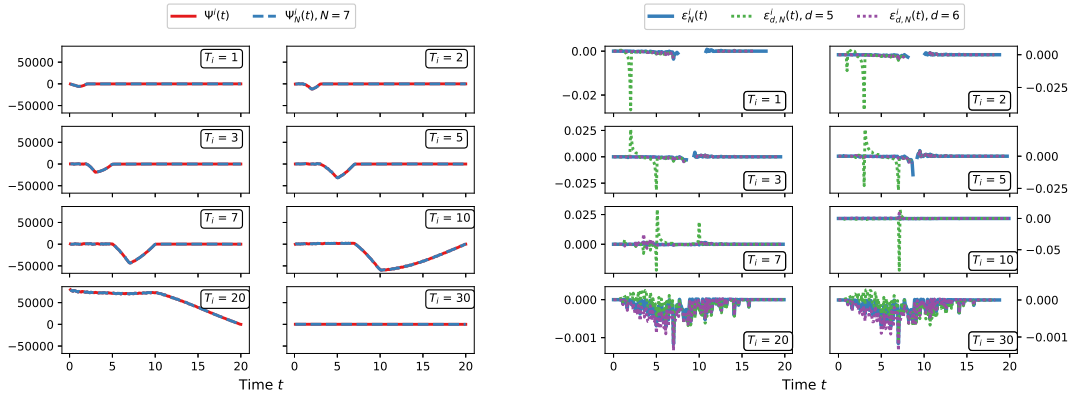


Figure 1: Left: Sensitivity profile obtained from exact approach $\Psi^i(t)$ and the full-order approximations $\Psi_N^i(t)$ for $1 \leq i \leq 8$. Right: Relative approximation errors of the full-order approach $\varepsilon_N^i(t)$, and of the low-order difference approach, $\varepsilon_{d,N}^i(t)$. $N = 7$, $d \in \{5, 6\}$. It should be noted that the relative approximation errors are negligible or confined to small peaks that have a minimal effect on the time integral of the sensitivity curve, which is ultimately utilized in xVA metrics.

$\bar{N} = 10000$ units of the currency. The fixed swap rate \bar{K} is chosen as the par rate $\bar{K} = 0.02226$, i.e., it holds $V(t_0) = 0$ at initial time $t_0 = 0$.

Initially, the number N of interpolation nodes needs to be determined. In practice, an approximation $g(t, \cdot)$ of the valuation function $V(t, \cdot)$ is already obtained during calculations of the expected exposure. This can be reused in the sensitivity approximations to the effect of improved numerical efficiency. For this experiment, we find that $N = 7$ interpolation nodes are sufficient to achieve a relative approximation error of the expected exposure below 1 bp,

$$\varepsilon_{EE} := \left| \max_{t \in [t_0, T]} \left(\frac{\widetilde{EE}(t_0, t) - EE(t_0, t)}{EE(t_0, t)} \right) \right| \approx 2.7 \times 10^{-5} < 1 \text{ bp}. \quad (6.2)$$

The output of the sensitivity calculations is a sensitivity profile of the expected exposure of the portfolio, which is a collection of all the sensitivities $\partial EE(t_0, t) / \partial K_i$, for certain K_i within the market instruments. In these numerical experiments, we consider sensitivity with respect to all market quotes used, $1 \leq i \leq 8$.

In the collection of graphs on the left of Figure 1, the sensitivity profile of the swap portfolio is shown. It is computed both with the full-order estimator of the sensitivities, $\Psi_N^i(t) \approx \frac{\partial \widetilde{EE}(t_0, t)}{\partial K_i}$, defined in Definition 4.1, and with the exact valuation approach $\Psi^i(t) \approx \frac{\partial EE(t_0, t)}{\partial K_i}$, given in (5.1). For each $i \in \{1, \dots, 8\}$, the market instruments K_i correspond with maturities $T_i \in \{1, 2, 5, 7, 10, 20, 30\}$. We observe that the magnitude of the sensitivities increases for $i \in \{1, \dots, 7\}$, up to the special quote K_7 , whose maturity is equal to the maturity of the swap in the portfolio itself, $T_7 = 20 = \bar{T}_m$. With respect to the market quote K_8 which is associated with the maturity $T_8 = 30$, there is nearly no sensitivity as the portfolio expires beforehand.

In the graphical sensitivity profiles, no difference between the full-order approach and the exact valuations is discernible. To demonstrate the differences, we continue with representations of the approximation errors, and additionally include the low-order difference approximation $\Psi_{d,N}^i(t)$, defined in Definition 4.2.

For each of the sensitivities with respect to K_i , $1 \leq i \leq 8$, we compute a relative approximation error of the full-order approach, and of the low-order difference

approximations. These relative errors are defined as the total approximation errors normalized with the exact result $\Psi_i(t)$, that is,

$$\varepsilon_N^i(t) := \frac{\Psi_N^i(t) - \Psi^i(t)}{\Psi^i(t)}, \quad \varepsilon_{d,N}^i(t) := \frac{\Psi_{d,N}^i(t) - \Psi^i(t)}{\Psi^i(t)}, \quad (6.3)$$

respectively. The relative errors are not well-defined for $\Psi^i(t) = 0$ and only numerically stable, when $\Psi^i(t)$ is not too close to zero. In these cases, we refrain from computing the relative error.

The approximation errors of the single-swap portfolio are given in the graphs to the right of Figure 1, computed for $N = 7$ and for $d = 5$, $d = 6$. At most times, all three considered methods display an extremely small relative approximation error below 0.1%. There are a few exceptions at individual time points, however, the error remains very small. For the low-order difference approaches, it can be found to be at most 7% for $d = 5$, 0.6% for $d = 6$, and 0.2% for the full-order approach ($d = N$). In the context of computing xVA functions, the expected exposures are often integrated over time, which further minimizes the impact of these spike-type errors. This is further investigated in Section 6.2.

The efficiency improvement provided by the method is highly promising. Recall that M is the number of Monte Carlo paths which can easily range in the tens of thousands for exposure simulations. Under the classical, exact valuation approach, it is necessary to compute M exact valuations at each time point for the unshocked market, and another $M \times n$ exact valuations for the n shocked markets considered. In this example, this attributes to 180000 exact valuations at each monitoring time t .

In the full-order approach with N interpolation nodes, at each time t only $(n+1) \times N$ exact valuations (63 in this example) are required. Additionally, the approximation function must be constructed, which attributes a numerical complexity of at worst $\mathcal{O}(N^2)$ per approximation function. Finally, the approximated portfolio valuations require M evaluations of the approximation function, which are always cheap to evaluate, regardless of the choice of the exact valuation function. For the low-order difference method, the number of exact valuations is further reduced to $N + n \times d$ (47 and 55, respectively, for $d = 5$, $d = 6$). This efficiency advantage increases further, as greater sensitivity profiles with more market shocks n are considered.

We demonstrate the advantage of using the low-order difference approach over the more straightforward full-order approach with a reduced degree N . To this end, we consider the low-order difference approach $\Psi_{d,N}^i(t)$ with $N = 7$ and $d = 5$, for which 47 exact valuations are required. It is compared to the full-order approaches $\Psi_N^i(t)$ for $N = 5$ and $N = 6$, with 45 and 54 exact valuations, respectively. This choice provides a similar number of exact valuations and we will demonstrate that the baseline parameter $N = 7$ in the low-order approach is not simply overspecified.

The findings are visualized in the collection of graphs on the left of Figure 2, where the relative errors of these three approaches are compared. From these error plots, we can observe that the low-order difference approximation performs substantially better than both alternative full-order approaches.

When a choice has to be made on how the number of interpolation nodes is allocated, this implies that it is advantageous to choose a larger value N at the cost of reducing the value d , rather than spreading the number of nodes uniformly to $d = N$ with a lower number of nodes N . An expected exposure calculation benefits from such a selection, since it utilizes only the unshocked portfolio approximation.

We conclude the experiments on the single-swap portfolio with a stressed market scenario, where the volatility η of the underlying interest rate $r(t)$ is now 5%,

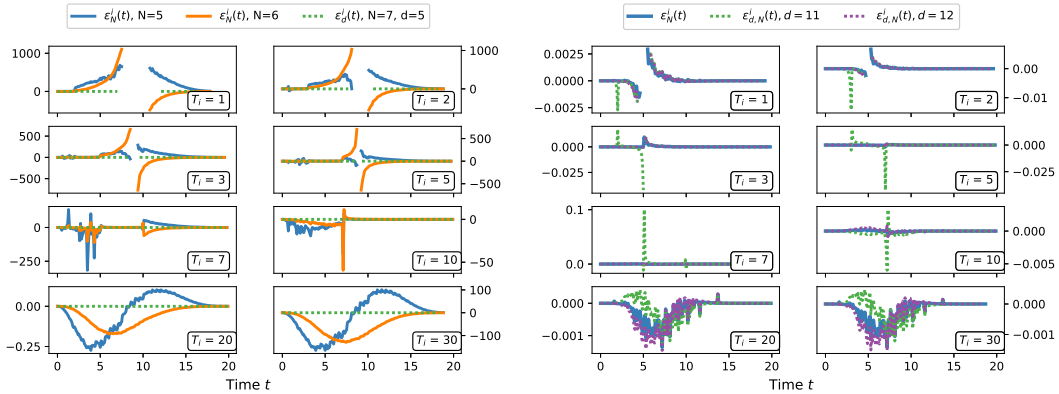


Figure 2: Left: The relative errors of the full-order approach $\varepsilon_N^i(t)$ obtained with fewer nodes $N = 5$ and $N = 6$, compared to the relative error of the low-order difference approach, $\varepsilon_{d,N}^i(t)$, for $N = 7$, $d = 5$. Note that the low-order difference approximation outperforms the full-order approximations with reduced values for N , even though they have similar numerical complexity.

Right: The relative approximation errors in a stressed scenario of high volatility $\eta = 5\%$ remain low. $N = 13$, $d \in \{11, 12\}$.

increased from previously 2%. To guarantee that the pricing error of the expected exposure adheres to the threshold, $\varepsilon_{EE} < 1$ bp, the number of interpolation nodes must be increased to $N = 13$. In the collection of graphs to the right of Figure 2, the corresponding relative errors are displayed for the full-order approach and the low-order difference approaches with 1 and 2 points removed, respectively. Magnitude and behaviour of the relative errors appear comparable to the observations made in the previous experiments with an unstressed volatility coefficient, which indicates a robustness of the method.

6.2. Large swap portfolio

Having examined the properties of the model in a simplified setting, we consider a more realistic portfolio of linear derivatives that involves a larger number of interest rate swaps. The exact composition of this portfolio is given in Appendix D. A choice of $N = 13$ interpolation nodes ensures an expected exposure estimation error close to the threshold of 1 basis point ($\varepsilon_{EE} \approx 1.2$ bp). However, we note that already with a smaller number of nodes $N = 9$, the maximal expected exposure error ε_{EE} is smaller than 7 basis points.

In Figure 3, we repeat the previously established experiment to obtain the relative sensitivity approximation errors. As in the previous experiments, we observe extremely small estimation errors. Particularly, it stands out that the largest relative estimation errors continue to be restricted to individual time points.

The expected exposures are not a risk metric in themselves, but rather a component of xVA computations. In many of these, the expected exposures are integrated over time, together with weight functions representing risks such as defaults or funding costs, see for example the credit valuation adjustment with wrong-way risk in (A.3). We introduce another error measurement that resembles the sensitivities of xVA computations in a general setting. This error, ζ_d^i , relates to the sensitivity of the




Figure 3: Relative approximation errors for the large swap portfolio remain small, analogous to the single swap portfolio in Figure 1. $N = 13$, $d \in \{11, 12\}$.

time-integrated expected exposures,

$$\left| \frac{\partial}{\partial K_i} \int_{t_0}^T \widetilde{\text{EE}}(t_0, t) dt - \frac{\partial}{\partial K_i} \int_{t_0}^T \text{EE}(t_0, t) dt \right| \leq \int_{t_0}^T \left| \frac{\partial}{\partial K_i} \widetilde{\text{EE}}(t_0, t) - \frac{\partial}{\partial K_i} \text{EE}(t_0, t) \right| dt \approx \int_{t_0}^T |\Psi_{d,N}^i(t) - \Psi^i(t)| dt =: \zeta_d^i. \quad (6.4)$$

We recall that $\Psi_{N,N}^i(t) = \Psi_N^i(t)$, so that ζ_d^i represents the error associated with the full-order approach for $d = N$, and the error associated with the low-order difference approach for $d < N$. To improve interpretability, we define the normalized versions of these error measurements by

$$\kappa_d^i := \frac{\zeta_d^i}{\int_{t_0}^T |\Psi^i(t_0, t)| dt}. \quad (6.5)$$

The findings are arranged in Table 1, where we tabulate the normalized errors κ_d^i for $1 \leq i \leq 8$ and a broad range of degrees for the low-order difference approximation, ranging from a linear difference approximation ($d = 2$) to the full-order approximation ($d = 13$). We observe that already with a low-order difference approach in $d = 7$ nodes, the normalized error is below 1% for all sensitivities. This translates to $N + d \times n = 69$ exact valuations, an improvement of over 40% compared to 117 valuations in the full-order approach. Both methods improve significantly on the $(n + 1) \times M$ evaluations in the classical, exact approach, which range in the tens to hundreds of thousands. In Table 2, an overview over the complexity improvements is provided.

7. Approximation and sensitivity of Bermudan swaptions

We now turn our attention to the approximation of exposure and sensitivity of Bermudan swaptions. These are non-linear interest rate derivatives composed of an underlying swap and the optionality of entering this swap at multiple exercise dates.

Table 1: The normalized, integrated expected exposure sensitivity approximation errors κ_d^i for the large swap portfolio. For each shock to K_i (associated with T_i), the lowest order d with an approximation error below 1% is highlighted. $N = 13$.

$d \backslash T_i$	1	2	3	5	7	10	20	30
2	1.3E+00	2.6E+00	2.6E+00	1.2E+00	1.0E+00	5.1E-01	1.4E-01	2.3E-01
3	1.2E-01	2.9E-01	2.5E-01	1.2E-01	1.3E-01	7.6E-02	1.2E-02	2.2E-02
4	1.0E-01	2.3E-01	3.6E-01	2.0E-01	1.8E-01	9.4E-02	1.3E-02	2.8E-02
5	2.0E-03	3.1E-03	1.3E-02	1.2E-02	1.6E-02	1.1E-02	1.8E-03	2.8E-03
6	1.4E-03	1.9E-03	1.0E-02	6.6E-03	1.2E-02	1.1E-02	2.2E-03	2.9E-03
7	4.3E-05	1.3E-04	1.3E-03	1.4E-03	2.2E-03	1.6E-03	3.0E-04	4.7E-04
8	2.2E-05	6.0E-05	7.5E-04	7.1E-04	1.2E-03	9.8E-04	1.4E-04	2.1E-04
9	8.0E-06	2.3E-05	2.6E-04	3.9E-04	7.9E-04	6.5E-04	8.8E-05	1.3E-04
10	1.0E-05	1.8E-05	2.0E-04	2.2E-04	6.1E-04	6.3E-04	1.5E-04	1.8E-04
11	4.9E-06	7.0E-06	7.0E-05	1.0E-04	2.8E-04	2.9E-04	9.3E-05	1.2E-04
12	5.7E-06	6.6E-06	1.1E-05	7.9E-06	3.3E-05	8.4E-05	4.5E-05	8.6E-05
13	3.9E-06	4.4E-06	7.4E-06	4.6E-06	2.6E-05	4.1E-05	4.3E-05	7.2E-05

Table 2: The number of exact valuations at each monitoring date t obtained from the classical approach and the approximation methods. We assume a Monte Carlo simulation with $M = 20000$ samples, $N = 13$ nodes in the unshocked market and $n = 8$ market shocks.

d	Low-order $N + dn$						Full-order $N + Nn$	Classical $M + Mn$
	7	8	9	10	11	12		
Number of exact valuations	69	77	85	93	101	109	117	180'000
Proportion of full-order approach	59%	66%	73%	79%	86%	93%	100%	> 1500%

This makes the valuation of a Bermudan swaption an optimal-exercise type problem, where at each exercise date, the (expected) pay-off from exercising the option (i.e., entering the underlying swap) has to be weighed against the continuation of the option to the next exercise date. This results in a high numerical valuation cost, particularly so in exposure simulations.

Efficient pricing algorithms have received much attention in the literature. In the seminal article of Longstaff and Schwartz (2001), the authors introduce the least squares Monte Carlo (LSMC) algorithm, which can be briefly described as a recursive, backwards search through the option values at the exercise times, where the continuation values are approximated by linear regressions. This allows for a simulation-based valuation, which only requires one set of paths per valuation time. Another notable approach has been introduced by Glau et al. (2021), where an iterative use of Chebyshev interpolation approximates the continuation values and allows for further complexity reduction in the exposure simulation of a Bermudan swaption.

In our outline of a possible approximation for the valuation of a Bermudan swaption $U(t, r(t))$ with underlying interest rate swap $V(t, r(t))$, we assume physical settlement, i.e. the Bermudan swaption is replaced by the underlying swap after exercise. Let S_1, \dots, S_L , $L \in \mathbb{N}$, be the exercise dates and let $\tau \in \mathbb{R} \cup \{\infty\}$ be the time of exercise, where $\tau = \infty$ indicates no exercise of the optionality. Then, the portfolio comprised

of this Bermudan swaption is given by

$$\Pi(t, r(t)) = \begin{cases} U(t, r(t)), & t < \tau, \\ V(r(t)), & t \geq \tau. \end{cases} \quad (7.1)$$

This combination of valuation functions introduces a new challenge to approximations. In an exposure simulation path $\omega \in \Omega$, the portfolio value at some point $r(t, \omega)$ may be either the value of the Bermudan swaption, $U(t, r(t, \omega))$ if $t < \tau(\omega)$, or the value of the underlying swap, $V(t, r(t, \omega))$ if the option has been previously exercised, i.e., $t \geq \tau(\omega)$.

This implies that there can be no single polynomial interpolation of the portfolio value. A possible resolution lies in the construction of two approximation functions $g_U(t, \cdot)$ and $g_V(t, \cdot)$, where g_U approximates the Bermudan swaption value U , and g_V the swap value V .

Then, the approximation of the portfolio valuation is given by

$$g(t, r(t, \omega)) = \begin{cases} g_U(t, r(t, \omega)), & t < \tau(\omega), \\ g_V(r(t, \omega)), & t \geq \tau(\omega), \end{cases} \quad (7.2)$$

which implies that the path-wise execution times $\tau(\omega)$ must be tracked in the exposure simulation. A numerically efficient way to estimate these execution times is an initial pass over all exposure simulation paths $r(t, \omega)$, for $t \in [t_0, S_n]$ with the LSMC algorithm, which allows for an approximation of the exercise time $\tau(\omega)$ of each path. Karlsson et al. (2016) remark that this approach constitutes an estimation of an exercise boundary. With such a boundary, it becomes possible to rephrase the exercise decision in terms of a threshold $r^*(t)$ in the underlying short rate, where the option is exercised at time S_k if $r(S_k) < r^*(S_k)$.

This implies that all inputs to the optionality component approximation $g_U(t, \cdot)$ come from a truncated distribution $r(t) | \{r(t) < r^*(t)\}$. The corresponding quadrature nodes can be computed by the Golub-Welsch algorithm (Golub and Welsch, 1969) from the moments of this truncated distribution.

For our numerical experiment, we consider a Bermudan swaption with exercise dates $S_k = k$ for $k \in \{1, \dots, 5\}$, strike rate \bar{K} and as an underlying the swap $V(t)$ described in Section 6.1. Whenever the Bermudan swaption is exercised, it transforms into a swap, as treated in the previous section. Therefore, we focus on monitoring dates prior to the expiry of the option, $0 \leq t \leq 5$, and solely on the approximation of the valuation function in unexercised paths, i.e., the approximation $g_U(t, \cdot)$ of the optionality component. For the sensitivity approximations, we restrict the profile to sensitivities with respect to $T_i \in \{1, 2, 3, 5\}$ years, where a large effect on the optionality component can be observed.

As indicated above, the quadrature nodes can be explicitly computed with the moments of the truncated distribution. In this experiment, the interest rate is simulated with the Hull-White dynamics given in (6.1), so the required moments are those of a truncated normal distribution which can be explicitly computed, see for example Burkardt (2014). In the left graph of Figure 4, we display the distribution of the samples $r(t, \omega)$ that are evaluated with the approximation of the optionality component, $g_U(t, \cdot)$, and we show the corresponding choice of quadrature interpolation nodes at time $t = 2$.

We find that with these quadrature nodes, a choice of $N = 15$ nodes yields a maximal expected exposure error $\varepsilon_{EE} < 3$ bp. In the graph on the right of Figure 4, we compare the sensitivity profile obtained from the LSMC algorithm with profile obtained

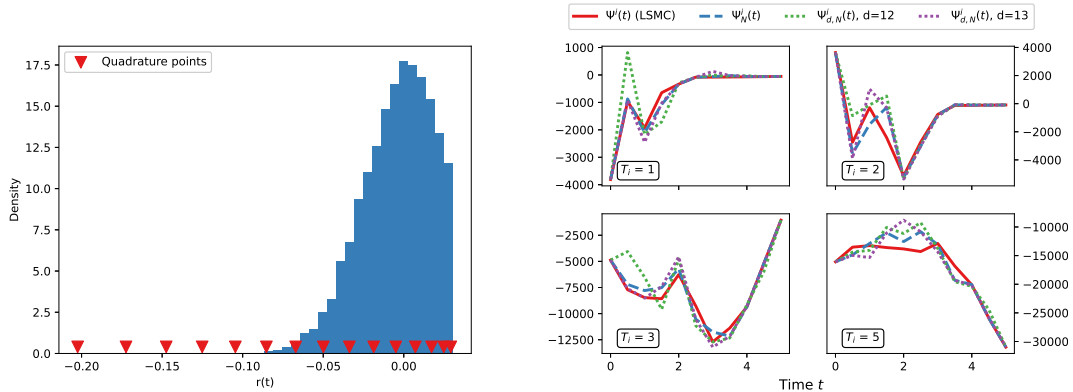


Figure 4: Left: Distribution of the risk factor in unexercised paths at time $t = 2$ and the corresponding quadrature nodes for $N = 15$. Right: Sensitivity profiles of the Bermudan Swaption, obtained from the full- and low-order approximation methods and the LSMC algorithm.

from the full-order and low-order approximation for $N = 15$, $d \in \{12, 13\}$. In this experiment, the high computational demand of computing expected exposures with the LSMC algorithm leads to relatively low accuracy of the ‘exact’ pricing approach, which makes it challenging to draw deep conclusions about the sensitivity approximation error. We recognize the sensitivity profiles computed in Figure 4 as plausible and emphasize the large numerical savings provided by the introduced methods.

8. Conclusion

We have shown that the computational demands of expected exposure sensitivity calculations can be greatly reduced in a setting of bump-and-revalue schemes, used to obtain xVA sensitivities with respect to the yield curve. We showed that the use of quadrature interpolation nodes allows for efficient approximations of the portfolio valuation functions, which is combined with a stochastic collocation technique to obtain the desired sensitivities. Two methods for the determination of the expected exposure sensitivity have been introduced. The first, full-order approach applies equal-degree polynomial approximations to shocked and unshocked valuation functions. The second, low-order difference approach offers an even stronger reduction in numerical complexity by utilizing a reduced order model to appraise the difference introduced to the valuation function by the market shock.

These methods can be implemented using basic numerical techniques, the computation of the quadrature nodes requires only matrix operations, and the estimators are obtained using classical polynomial interpolation algorithms.

The convergence of the low-order difference approximation to the full-order approximation of the sensitivity, and of the full-order approximation of the sensitivity to the classical, exact-valuation-based sensitivity have been studied.

Numerical experiments have demonstrated the accuracy and numerical efficiency of the method. Further, experiments indicate that the low-order difference approach, where an increased degree approximation of the unshocked portfolio is combined with a difference approximation of a lower degree, may be used to improve the accuracy for both the expected exposure simulation and the sensitivity computations.

Applications of approximation techniques in exposure simulations have recently provided fruitful research topics, but further challenges and opportunities remain. The

approach to sensitivity computations proposed in this article is based on an interpolation node selection which is entirely domain-driven, that is, the selection of nodes only depends on the distribution of the underlying risk factors. Beyond a certain regularity to ensure the approximability of the portfolio valuation functions, little is assumed about their functional structure. Adaptive algorithms that infer structures such as domains of functional variation from pointwise evaluations may provide a useful improvement in efficiency and complexity, whilst preserving the advantage of not requiring strong assumptions about the composition of the portfolios considered.

Acknowledgments

The authors thank an anonymous referee for valuable comments that helped improve this article. This research is part of the ABC–EU–XVA project and has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska–Curie grant agreement No. 813261.

References

- M. Abramowitz and I. A. Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1964.
- J.-P. Berrut and L. N. Trefethen. Barycentric Lagrange interpolation. *SIAM review*, 46(3):501–517, 2004. doi: 10.1137/S0036144502417715.
- D. Brigo and F. Mercurio. *Interest rate models: theory and practice: with smile, inflation, and credit*. Springer, Berlin New York, 2006. doi: 10.1007/978-3-540-34604-3.
- R. L. Burden, J. D. Faires, and A. M. Burden. *Numerical analysis*. Cengage learning, 2015. doi: 10.13140/2.1.4830.2406.
- J. Burkardt. The truncated normal distribution. *Department of Scientific Computing Website, Florida State University*, 1:35, 2014.
- L. Capriotti, Y. Jiang, and A. Macrina. AAD and least-square Monte Carlo: Fast Bermudan-style options and XVA Greeks. *Algorithmic Finance*, 6(1-2):35–49, Oct. 2017. doi: 10.3233/af-170201.
- M. Gaß, K. Glau, M. Mahlstedt, and M. Mair. Chebyshev interpolation for parametric option pricing. *Finance and Stochastics*, 22(3):701–731, Apr. 2018. doi: 10.1007/s00780-018-0361-y.
- K. Glau, P. Herold, D. B. Madan, and C. Pötz. The Chebyshev method for the implied volatility. *Journal of Computational Finance*, pages 1–31, Dec. 2019a. doi: 10.21314/jcf.2019.375.
- K. Glau, M. Mahlstedt, and C. Pötz. A new approach for American option pricing: The dynamic Chebyshev method. *SIAM Journal on Scientific Computing*, 41(1):B153–B180, Jan. 2019b. doi: 10.1137/18m1193001.
- K. Glau, R. Pachon, and C. Pötz. Speed-up credit exposure calculations for pricing and risk management. *Quantitative Finance*, 21(3):481–499, 2021. doi: 10.1080/14697688.2020.1781236.
- A. Gnoatto, A. Picarelli, and C. Reisinger. Deep xVA solver – a neural network based counterparty credit risk management framework, 2020. URL <https://arxiv.org/abs/2005.02633>.
- G. H. Golub and J. H. Welsch. Calculation of Gauss quadrature rules. *Mathematics of computation*, 23(106):221–230, 1969.
- A. Green. *XVA: Credit, Funding and Capital Valuation Adjustments*. John Wiley & Sons, 2015. doi: 10.1002/9781119161233.
- J. Gregory. *The xVA Challenge: Counterparty Risk, Funding, Collateral, Capital and Initial Margin*. John Wiley & Sons, 2020. doi: 10.1002/9781119508991.
- L. A. Grzelak. Sparse grid method for highly efficient computation of exposures for xVA. *Applied Mathematics and Computation*, 434:127446, 2022. doi: 10.1016/j.amc.2022.127446.
- L. A. Grzelak, J. Witteveen, M. Suárez-Taboada, and C. Oosterlee. The stochastic collocation Monte Carlo sampler: highly efficient sampling from expensive distributions. *Quantitative Finance*, 19: 1–18, 06 2018. doi: 10.1080/14697688.2018.1459807.
- P. S. Hagan and G. West. Interpolation methods for curve construction. *Applied Mathematical Finance*, 13(2):89–129, 2006. doi: 10.1080/13504860500396032.
- B. N. Hugel and A. Savine. LSM reloaded - differentiate xVA on your iPad mini. *SSRN Electronic Journal*, 2017. doi: 10.2139/ssrn.2966155.
- J. Hull and A. White. Pricing interest-rate-derivative securities. *The Review of Financial Studies*, 3 (4):573–592, 1990. doi: 10.1093/rfs/3.4.573.

- P. Karlsson, S. Jain, and C. W. Oosterlee. Fast and accurate exercise policies for Bermudan swaptions in the LIBOR market model. *International Journal of Financial Engineering*, 03(01):1650005, Mar. 2016. doi: 10.1142/s2424786316500055.
- F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: a simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147, 2001. doi: 10.1093/rfs/14.1.113.
- C. W. Oosterlee and L. A. Grzelak. *Mathematical Modeling and Computation in Finance: With Exercises and Python and Matlab Computer Codes*. World Scientific Publishing Company, 2019. doi: 10.1142/q0236.
- L. N. Trefethen. *Approximation Theory and Approximation Practice, Extended Edition*. SIAM, 2019. doi: 10.1137/1.9781611975949.
- D. Xiu and J. S. Hesthaven. High-order collocation methods for differential equations with random inputs. *SIAM Journal on Scientific Computing*, 27(3):1118–1139, 2005. doi: 10.1137/040615201.

Appendix A. Extension to more general xVA models on the example of wrong-way risk

In the following, we demonstrate how portfolio valuation approximation can be applied to more complex xVA models, by example of the unilateral CVA with wrong-way risk (WWR).

Generally, the CVA is constructed from three components. The *probability of default* of the counterparty (PD), the *expected (positive) exposure* (EE) and the *loss given default* (LGD), which describes what percentage of the exposure is lost in a default event. A common and simplifying CVA approximation is based on the assumption of independence between exposure and the probability of default (see, for example, Gregory (2020)). The resulting credit valuation adjustment is given by

$$\text{CVA}_1(t_0) := \text{LGD} \cdot \sum_{k=1}^R \text{EE}(t_0, t_k) \text{PD}(t_{k-1}, t_k), \quad (\text{A.1})$$

where $t_0 < \dots < t_R$, $R \in \mathbb{N}$, is a discretization of the time horizon. The independence assumption which allows for factorization of EE and PD is not always warranted, as the exposure may rely on the same risk factors which influence the default risk. Depending on the sign of correlation, this is known as wrong-way risk (when exposure increases together with the probability of default) or right-way risk (when exposure decreases as the probability of default rises). Let τ denote the time of counterparty default and assume that at the time of valuation t_0 , the default has not yet occurred, $\tau > t_0$. Then, a more general formulation of the CVA is given by

$$\text{CVA}_2(t_0) := \text{LGD} \cdot \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\frac{B(t_0)}{B(\tau)} \mathbb{1}_{\{\tau \leq T\}} V^+(\tau) \right]. \quad (\text{A.2})$$

One way to introduce correlations is through stochastic hazard rates $y(t)$, which are continuous stochastic processes that describe the intensity default process. Correlating these hazard rates with the underlying is one way to induce an analytically solvable dependency structure between default times and exposure. Analogous to the stochastic intensity model in Chapter 22.7 of Brigo and Mercurio (2006), it is possible to show that the previous equation can be rewritten as

$$\text{CVA}_2(t_0) = \text{LGD} \cdot \int_{t_0}^T \mathbb{E}_{t_0}^{\mathbb{Q}} \left[\frac{B(t_0)}{B(s)} \left(e^{-\int_{t_0}^s y_u du} y_s \right) V^+(s) \right] ds. \quad (\text{A.3})$$

This expected value contains an additional factor corresponding to the probability of default, $D(s) := \exp(-\int_{t_0}^s y_u du) y_s$, which is governed by the stochastic hazard rates.

The asset valuation, however, is not impeded. Therefore, a Monte Carlo experiment analogous to (2.2) can be conducted, where an additional risk factor $y(t)$ is simulated. The approximation techniques discussed in this article can be applied in the usual way to replace the exact portfolio valuations V with a polynomial approximation function g , to the effect of the same efficiency improvements as before.

Appendix B. Explicit discount factor sensitivity

We consider the derivative of the discount factor

$$\frac{\partial}{\partial K_i} \frac{B(t_0; \omega)}{B(t; \omega)} = \frac{\partial}{\partial K_i} \exp\left(-\int_{t_0}^t r(s; \omega) ds\right). \quad (\text{B.1})$$

In the following, we will consider the explicit choice of the 1-factor Hull–White model for the interest rate $r(t)$, but the methodology can be extended to many other exogenous interest rate models (see Brigo and Mercurio (2006) for an overview of such models). Then, the short rate $r(t)$ can be explicitly expressed in terms of a stochastic process $u(t)$, and a purely deterministic function $\psi(t)$, such that at any time it holds $r(t; \omega) = u(t; \omega) + \psi(t)$. The deterministic function $\psi(t)$ imposes the term structure of interest rates, i.e. the initial yield curve, onto the model, whereas the stochastic process $u(t)$ only depends on the yield curve through its parameter calibration procedure. As explained in Chapter 21 of Green (2015), the remaining model parameters are not usually recalibrated when sensitivities to the yield curve are computed. Under this assumption, it holds that

$$\begin{aligned} \frac{\partial \exp(-\int_{t_0}^t r(s; \omega) ds)}{\partial K_i} &= \exp\left(\int_{t_0}^t -u(s; \omega) ds\right) \frac{\partial \exp(\int_{t_0}^t -\psi(s) ds)}{\partial K_i} \\ &= \exp\left(\int_{t_0}^t -u(s; \omega) ds\right) \exp\left(\int_{t_0}^t -\psi(s) ds\right) \left(-\int_{t_0}^t \frac{\partial \psi(s)}{\partial K_i} ds\right). \end{aligned} \quad (\text{B.2})$$

The deterministic function $\psi(t)$ is explicitly known, in this example of the 1-factor Hull–White model it is given for $t_0 = 0$ by

$$\psi(t) = f^r(0, t) + \exp(-\lambda t) \frac{\eta^2}{\lambda^2} (\cosh(\lambda t) - 1), \quad (\text{B.3})$$

where λ is the speed of mean reversion parameter, η is the volatility parameter, and $f^r(0, t)$ is the instantaneous forward rate associated with the model. Under the yield curve, respectively its associated zero-coupon bond curve $P(0, \cdot)$ given in (3.2), the instantaneous forward rate at time t equals

$$f^r(0, t) = \frac{-\partial \log(P(0, t))}{\partial t}. \quad (\text{B.4})$$

Since the model parameters λ, η are not recalibrated to the change in the yield curve, it can be shown that it holds

$$\frac{\partial \psi(t)}{\partial K_i} = \frac{\partial f^r(t)}{\partial K_i} = \frac{\partial \left(-\frac{\partial \log(P(0, t))}{\partial t}\right)}{\partial K_i} = -\frac{\partial}{\partial t} \left(\frac{1}{P(0, t)} \frac{\partial P(0, t)}{\partial K_i}\right). \quad (\text{B.5})$$

The derivative with respect to the market quote, $\partial P(0, t)/\partial K_i$ is typically not analytically available since the zero-coupon bond curve $P(0, t)$ is obtained from the market quotes by means of a curve-fitting scheme.

Appendix C. Sensitivity in the yield curve construction

As noted below (B.5), the derivatives of the zero-coupon bonds with respect to the market quote K_i are not generally available, since these bonds are often obtained through contrived curve-fitting schemes. In generality, each zero-coupon bond $P(0, t)$ depends on all of the market instruments $\varphi_1, \dots, \varphi_n$ (given in Definition 3.2) on which the yield curve is based. We may emphasize this by writing

$$P_t(\varphi_1, \dots, \varphi_n) := P(0, t). \quad (\text{C.1})$$

This implies that the inner derivative in (B.5) is expressed by

$$\frac{\partial P(0, t)}{\partial K_i} = \frac{\partial P_t(\varphi_1, \dots, \varphi_n)}{\partial K_i} = \sum_{j=1}^n \frac{\partial P_t(\varphi_1, \dots, \varphi_n)}{\partial \varphi_j} \frac{\partial \varphi_j}{\partial K_i}. \quad (\text{C.2})$$

Importantly, the effect of each market instrument on each zero-coupon bond is model-dependent. For example, when the yield curve is constructed with the Newton-Raphson method, a Jacobian of the form

$$J_\varphi = \begin{pmatrix} \frac{\partial \varphi_1}{\partial P(0, S_1)} & \cdots & \frac{\partial \varphi_1}{\partial P(0, S_n)} \\ \vdots & \ddots & \vdots \\ \frac{\partial \varphi_n}{\partial P(0, S_1)} & \cdots & \frac{\partial \varphi_n}{\partial P(0, S_n)} \end{pmatrix} \quad (\text{C.3})$$

is obtained, where each maturity S_j is a spine point associated with exactly one market instrument φ_j , $j \in \{1, \dots, n\}$. Under some restrictions on the choice of market instruments such that the inverse function theorem holds for J_φ , the inverse Jacobian

$$J_P = \begin{pmatrix} \frac{\partial P(0, S_1)}{\partial \varphi_1} & \cdots & \frac{\partial P(0, S_1)}{\partial \varphi_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial P(0, S_n)}{\partial \varphi_1} & \cdots & \frac{\partial P(0, S_n)}{\partial \varphi_n} \end{pmatrix} \quad (\text{C.4})$$

can be found explicitly with the relation $J_P = J_\varphi^{-1}$. It can not be guaranteed that every maturity t , for which a zero-coupon bond $P(0, t)$ is required, is also an element of the spine points $\{S_1, \dots, S_n\}$. This brings about the need for another approximation. Direct interpolation between the interpolation points $\{\frac{\partial P(0, S_\ell)}{\partial \varphi_j} : \ell \in \{1, \dots, n\}\}$ often does not admit good results. In these cases, it is preferable to deploy another chain rule,

$$\frac{\partial P(0, t)}{\partial \varphi_j} = \sum_{\ell=1}^n \frac{\partial P(0, t)}{\partial P(0, S_\ell)} \frac{\partial P(0, S_\ell)}{\partial \varphi_j}, \quad (\text{C.5})$$

and interpolate between the missing terms $\partial P(0, t)/\partial P(0, S_\ell)$ instead. Clearly, even in such a semi-analytical approach, the yield curve construction methodology effects the sensitivities. The same holds in direct “bump-and-revalue” schemes, as used within this article, directly through the construction of the shocked and unshocked yield curves.

Appendix D. Parameters of the numerical experiments

In the numerical experiments, a yield curve is constructed based on (artificial) market quotes for interest rate swaps, given in Table D.3. The composition of the large swap portfolio is specified in Table D.4.

Table D.3: Market instruments (interest rate swaps) used in the yield curve construction.

	i :	1	2	3	4	5	6	7	8
Maturity T_i :		1	2	3	5	7	10	20	30
Swap rate K_i (in %):		0.04	0.16	0.31	0.81	1.28	1.62	2.22	2.30

Table D.4: Specification of the interest rate swaps comprising the large swap portfolio.

Sign	Notional	Fixed Rate	Maturity	Start Date	Payments per Year
-1	10000	0.022	20	0	2.0
-1	8333	0.042	20	0	2.0
-1	8333	0.042	21	0	1.9
1	8333	0.042	24	0	1.7
1	8333	0.042	17	0	2.4
1	8333	0.042	26	0	1.5
1	8333	0.042	19	5	2.9
1	8333	0.042	40	10	1.3
-1	8333	0.042	19	3	2.5
-1	8333	0.042	20	7	3.1
1	8333	0.042	20	2	2.2
-1	8333	0.042	20	0	2.0
-1	8333	0.042	20	0	2.0