Computing Valuation Adjustments for Carbon Forward Contracts: A Market Model Based Approach

Master Thesis Msc Quantitative Finance (FEM61008)

Thomas Geuzinge (493105)

Supervisor: dr. Evgenii Vladimirov
Second Assessor: prof. dr. Michel van der Wel

Date Final Version: February 27, 2024

Abstract

This paper presents a novel approach to compute Credit Valuation Adjustments and Debt Valuation Adjustments for carbon forward contracts by simulating forward interest rates and forward exchange rates using the Commodity Interest Rate Market (CIRM) model. Existing market models in the literature have not been applied to the market for European CO_2 emission allowances. At the same time, Basel III regulations require that institutions take counterparty credit risk into account in the fair value of their derivative contracts. We calibrate the CIRM model using market-implied volatilities and market-observable interest and exchange rates. With the calibrated CIRM model we simulate forward interest rates and forward exchange rates, and our findings indicate that Valuation Adjustments can be accurately calculated based on these simulated rates. Next to this, we observe that the behavior of Valuation Adjustments follows economic rationale when manually adjusting the volatility and correlation parameters of the CIRM model. Therefore, we believe that the proposed methodology can be of interest to practitioners in the financial industry.

Erasmus University Rotterdam
Erasmus School of Economics

Contents

1 Introduction

Many large corporations have been concerned with reducing their carbon footprints over the last few years. One of the difficulties companies face with shrinking their carbon footprint is the uncertainty of future carbon emission allowance prices. Carbon emission allowances are traded on the market but their prices depend both on the supply of allowances (which is highly influenced by government policy) as well as the demand for allowances. To prevent this uncertainty in allowance prices, financial institutions can offer forward derivative contracts regarding CO_2 emission allowances (henceforth named "carbon forwards") to their corporate customers. However, as with any type of derivative contract, offering a derivative contract to customers introduces risk for an institution. Two important types of risk are counterparty credit risk and own credit risk. To take these two risk types into account in the fair value of the derivative contract, an institution can calculate Valuation Adjustments (XVA). The valuation adjustment for counterparty credit risk is also known as the Credit Valuation Adjustment (CVA), while the adjustment for own credit risk is called the Debt Valuation Adjustment (DVA). XVA are typically calculated using Monte Carlo simulation of diffusion-type stochastic market models. However, XVA for carbon forward contracts have not been researched in the current academic literature. Therefore, to take the effects of credit risk on the fair value of carbon forward contracts more accurately into account, new research is required.

Carbon emission allowances are traded on the European Union Emissions Trading System (EU ETS) market, the largest market for CO_2 emission allowances in the world. The market is a direct consequence of the Kyoto Protocol that was formulated in 1997 to reduce the emission of greenhouse gasses. The carbon market in the EU is a so-called 'cap and trade' market, which means that a cap on the supply of emission certificates is set each year and that the available certificates can be exchanged between polluting parties. By the end of the year, each company should have bought enough allowances to cover their yearly emissions or they face significant fines. Also, the cap on the maximum amount of allowances in the market is reduced each year such that the EU reaches its goal of becoming climate-neutral by 2050 (European Commission, 2023). The EU ETS system has reduced the emissions from power and industrial companies by 37% since 2005 (European Commission, 2023). Additionally, the size of carbon markets is steadily increasing, with global transactions amounting to roughly 10 billion USD in 2010 to 96 billion USD in 2022 (World Bank, 2023). The market size is primarily growing because of the price increase of allowances, thus making financial instruments that fix the future allowance price more attractive for large and polluting corporations.

Accurate estimation of the fair value of carbon forwards has numerous benefits, especially for the seller (e.g. banks and/or other financial institutions). Since the release of the Basel III regulatory framework in 2010, institutions are required to adjust the prices of financial derivatives on their balance sheets based on counterparty credit risk (Härle et al., 2010). They can take this credit risk into account by adding or subtracting XVA from the values of their contracts. Moreover, institutions charge (a part of) their XVA to their customers if the XVA exceed pre-determined thresholds. This way, they reduce their risk exposure related to each of their clients. Preferably, the XVA are not too high as this will be unattractive to customers, but the XVA should also not be too low as this leaves an institution with an open risk position as the risk is underpriced. Accurate computation of XVA is thus of great importance for both the balance sheet of the seller as well as from a compliance and risk viewpoint. Buyers of carbon forwards (often large corporations in the manufacturing or energy industries), in turn, also benefit from correctly calculating the fair value of carbon forwards. This is because correct estimation gives them more certainty regarding their future costs concerning carbon emission allowances when buying carbon forwards instead of buying allowances directly on the market. With this additional certainty on future costs, companies can make more informative decisions on whether they should buy carbon emission allowances for the foreseeable future or implement a strategy that reduces their carbon footprint. Furthermore, with carbon forwards companies can pay for their allowances at a later point in time than if they were to finance them immediately, which allows companies to use the money for other purposes in the meantime.

The academic relevance of this research is multi-faceted. Some researchers have focused on the market for carbon emission allowances, but most of these scholars focused on modeling the carbon allowance price (Chang & Wang, 2015; Daskalakis et al., 2009; Seifert et al., 2008). While this is an important topic in the carbon market, modeling only the carbon spot price is not sufficient to compute XVA, which is the goal of this paper.

To simulate XVA, one can use a market model for (forward) interest and exchange rates. Literature on market models is extensive, as market models describe the dynamics of interest rates which, in turn, are used in a large number of applications. So far, previous research into market models by e.g. Pilz and Schlögl (2013), Karlsson, Pilz, and Schlögl (2016) and Fernández et al. (2015) have calibrated LIBOR Market Models for interest rates to several types of commodities and commodity contracts. This research aims to extend this trend by calibrating a modified version of the LIBOR Market Model (which we will term the Commodity Interest Rate Market (CIRM) model) to carbon emission allowance price data. Since carbon emission allowances are only becoming more important in the future due to the current focus on combating climate change, this could be a relevant addition to the existing literature.

Additionally, this research will extend the current academic literature regarding the computation of XVA. Quite some papers on XVA simulation use the Least-Squares Monte Carlo (LSMC) algorithm of Longstaff and Schwartz (2001) as a basis for the XVA calculation. This algorithm has its limitations, which are partially overcome by extensions proposed by Areal et al. (2008) and Gamba (2003). However, originally the LSMC algorithm is developed to model the future value of a derivatives portfolio, not to determine XVA. Fortunately, Joshi and Kwon (2016) bridge this gap in the literature, but they evaluate interest rate swap contracts instead of carbon forwards. Note that some research has been performed into CVA for commodity swaps and forwards (e.g. Baldeaux and Platen (2015), Brigo et al. (2009), and Green (2015)), but these papers do not cover the carbon market and also do not incorporate a sophisticated market model in their research. Therefore, this paper is the first to combine a (modified) Commodity LIBOR Market Model with XVA for carbon forward contracts.

The previously described gaps in the literature combined with the practical importance of accurately calculating XVA has led us to formulate the following research question: How can Valuation Adjustments for carbon forwards be calculated in the framework of a Commodity Interest Rate Market model? To this end, we construct and calibrate a Commodity Interest Rate Market model that simulates forward exchange rates and forward interest rates. Then, we use these rates to compute XVA for netted and possibly collateralized derivative portfolios that include carbon forward contracts. Roughly speaking, a CVA is calculated by taking the product of a counterparty's average Expected Positive Exposure, its Loss Given Default and its Probability to Default. The DVA is determined similarly, but then by using the average Expected Negative Exposure instead of the average Expected Positive Exposure.

To investigate whether the XVA results that are obtained based on the CIRM model are valid, we test the model in real-life applications. In some of these applications, we compare the CIRM model-based results to the results based on an analytical benchmark XVA model. This benchmark model has an analytical approximation for the EPE and ENE and uses market-observable input parameters. Moreover, we check the correctness of the computed XVA by varying the input parameters of the CIRM model and using economic reasoning to see whether the results adjust as expected when we change the model parameters. Lastly, we assess whether the CIRM model can compute XVA for a portfolio of derivatives as well.

Our results show that the CIRM model simulates forward interest rate and forward exchange rate curves that closely match the market-observed curves. We observe some small differences between the curves, but after further investigation we find that they are not statistically significant. Furthermore, when allowing for more simulation paths in the CIRM model, we see that the differences diminish. However, this comes at the cost of computation time. The comparison of the exposure profiles based on the CIRM and the benchmark model shows that no clear distinction between the two models can be made. Both models give almost identical exposure profiles for a carbon forward over time, thus leading to similar XVA results as the PD and LGD are the same for both models. When we check the validity of the CIRM model using economic rationale, we see that manually increasing the carbon volatility parameter in the CIRM model leads, as expected, to increased exposure profiles. An explanation for this is that more volatile carbon exchange rates attain higher (and lower) levels. The exchange rates, in turn, have a substantial influence on the exposure levels. Additionally, we find that the CIRM model can also cope with collateralized carbon forward contracts and with a portfolio of derivative contracts. Finally, we assess whether the CIRM model produces reasonable XVA for (portfolios of) other derivative contracts and we conclude that the CIRM model is indeed able to do so.

The remainder of this paper is organized as follows. In Section 2, we discuss the relevant academic literature on the carbon market, interest rate modeling, and valuation adjustments. Section 3 describes the Commodity Interest Rate Market model that will be implemented in this research. Then, in Section 4, we disclose how we are going to calibrate the market model. Section 5 shows how the valuation adjustments are going to be determined. Then, in Section 6 we discuss the findings of this paper and, lastly, Section 7 summarizes the main findings and discusses the limitations of this research.

2 Literature Review

This section describes the existing academic literature on the most important concepts of this paper. First, we give a brief overview of the literature related to the carbon emission allowance market. Second, we examine the literature regarding interest rate modeling. Lastly, we touch upon the literature regarding Valuation Adjustments (XVA) and their computation.

2.1 Carbon Emission Allowance Market

Previous research into the carbon emission allowance market has primarily focused on modeling the dynamics of the carbon spot price. This can be done in numerous ways, e.g. Benz and Trück (2009) use a regime-switching GARCH model that can also be used to forecast the future carbon price, while Lovcha et al. (2022) employ a Structural Vector Autoregression model that links prices from energy sources to the carbon spot price. Seifert et al. (2008) present a stochastic equilibrium model for the carbon spot price and find that CO_2 price processes should allow for time- and price-dependent volatility. While these papers are all important for traders and risk managers from a trading and hedging perspective, none of these papers show how a carbon derivative contract can be evaluated within their frameworks. Daskalakis et al. (2009) develop a framework with a continuous-time jump-diffusion process for the spot price that can be utilized to value carbon futures and forward contracts. Similarly, Chang and Wang (2015) derive partial differential equations for carbon derivative prices that allow for stochastic carbon-related yields in a no-arbitrage framework.

However, modeling the option value of a carbon forward is only one step in the XVA calculation of a portfolio that consists of carbon forward contracts. Brigo et al. (2009) nicely show how CVA can be computed with a CIR model for a single commodity swap, but this approach cannot be used in a setting with multiple trades. Alternatively, one could skip the difficult simulation of interest rates and directly compute the value of a derivatives portfolio by using the Black-Scholes (Black & Scholes, 1973), Black (Black, 1976), or Garman-Kohlhagen (Garman & Kohlhagen, 1983) formulas. While this non-simulated approach gives a decent approximation of the XVA, this methodology does not work well for options with a forward premium (which arises when the expected future price is higher than the spot price). Therefore, the simulated approach that uses a comprehensive interest rate market model is preferable over the previously discussed methods for calculating XVA for a derivatives portfolio including carbon forwards.

2.2 Interest Rate Modeling

Interest rate modeling has been a topic of academic interest ever since the seminal work of Vasicek (1977). This first term structure model is still relatively simplistic and focuses on modeling the dynamics of short (interest) rates, which is the instantaneous borrowing and lending rate (Vasicek, 1977). Its simplicity follows from the fact that it assumes that market risk (also called systematic risk) is the only driving factor behind the dynamics of interest rates. Also, it assumes that drift and diffusion parameters are constant over time. The Vasicek model postulates that the entire interest rate curve can be modeled based on a stochastic differential equation for the dynamics of the short rate. The interest rate curve allows practitioners to calculate bond prices, as these prices are exponentially affine functions of interest rates. However, due to the many restrictive assumptions of the Vasicek model, these bond prices were not in accordance with the empirically observed bond prices in the market.

To better match the observed bond prices, several improved interest rate models were developed such as the CIR model (Cox et al., 1985), the Ho-Lee model (Ho & Lee, 1986), the Hull-White model (Hull & White, 1990), the HJM framework for forward rates (Heath et al., 1992) and the BGM model (Brace et al., 1997). The CIR model also assumes that market risk is the only underlying factor for the short rate, but in contrast to the Vasicek model, it states that the volatility of the short rate is time-varying and depends on the value of the short rate itself. Nonetheless, this model is also still quite restrictive and does not allow for negative interest rates, which is not in line with empirical observations on (forward) interest rates (Rostagno et al., 2021). The Ho-Lee and Hull-White models for the short rate can produce negative interest rates, and they differ from the Vasicek model in the fact that they allow for a time-varying mean level of the short rate. This makes it easier for these models to match the observed interest rates and bond prices in the market, but due to the one-factor nature of both models, they struggle to capture complex term structure dynamics like e.g. a volatility smile.

A different approach to determining the interest rate curve is to focus on modeling (instantaneous) forward rates instead of short rates. Instantaneous forward rates are rates that cover an infinitesimally small period in the future. This approach is taken in the HJM framework of Heath et al. (1992). The HJM framework tries to capture the entire forward rate curve by describing the dynamics for instantaneous forward rates, which can then be combined with an initial curve to simulate the forward rate curve. Based on this curve, bond prices can be determined by expressing them as functions of the forward rates. However, because the model describes instantaneous forward rates, the model tends to an infinite number of dimensions which makes it computationally intensive. A solution to this problem is provided by the BGM model, which describes the dynamics of simple forward rates. Since the model originally focused on London Interbank Offered Rate (LIBOR) forward rates, it is also called the LIBOR Market Model (LMM) after contributions were made to the model by Miltersen et al. (1997) and Jamshidian (1997).

The LMM is of particular interest, as it is the precursor of the market model that is implemented in this paper. It is called a market model because it differs from the aforementioned models in the sense that the LMM describes market-observable forward rates directly, instead of e.g. unobservable short rates as in the Vasicek model. Since the forward rates are observable in the market, calibration of the LMM is substantially easier than the calibration of the previously described models. The LMM is an arbitrage-free model that assumes that forward rates are log-normally distributed. Additionally, when the interest rate follows the dynamics described in the LMM, the Black 76 formula (Black, 1976) holds for interest rate options. This allows the volatility parameters in the differential equations of the LMM to be calibrated using the Black formula combined with observed market prices on interest rate options.

Yet, also the LMM has its disadvantages. For example, due to the log-normality assumption on the forward rates, the LMM forward rates cannot be negative. Recent empirical observations actually show that interest and forward rates can be negative (Rostagno et al., 2021), thus an interest rate model should allow for this behavior. A solution to this problem is given in Brigo and Mercurio (2006), where a displacement term is added to the forward rates to allow for negative forward rates. Additionally, a more practical problem of the LMM is that it only describes the dynamics of forward rates for one single currency. In reality, multiple currencies and their forward rates are often of importance. This is the case in this paper, as we will evaluate a derivatives portfolio that generates cash flows in more than one currency. To value this portfolio using just one market model for all interest rates, the Multicurrency LIBOR Market Model (MCLMM) was introduced by Schlögl (2002). Here, the forward rates of different currencies can be linked to one another via forward exchange rates. To use an MCLMM in the context of pricing commodity derivatives, Pilz and Schlögl (2013) note that a commodity market can be seen as a currency market. Based on this observation, they extend the MCLMM to the Commodity LIBOR Market Model (CLMM). Since carbon emission allowances can be seen as a commodity (Benz & Trück, 2009), we select the CLMM to be the cornerstone of the Commodity Interest Rate Market model (CIRM model), the market model that is developed in this paper.

2.3 Valuation Adjustments

Since the Basel III regulatory changes following the global financial crisis that started in 2008, modeling XVA has become a more popular topic in academic literature. Many different valuation adjustments exist, such as the Credit Valuation Adjustment (CVA), Debt Valuation Adjustment (DVA), Funding Valuation Adjustment (FVA), Capital Valuation Adjustment (KVA) and Margin Valuation Adjustment (MVA) (Green, 2015). The considered valuation adjustments in this paper are the adjustments related to default risk, i.e. the CVA and the DVA. Therefore, this literature section will focus only on the literature regarding these two valuation adjustments. The CVA consists of three components, the Expected Positive Exposure (EPE), the Probability to Default (PD), and the Loss Given Default (LGD). Similarly, the DVA can be decomposed in the PD, LGD, and the Expected Negative Exposure (ENE).

Oftentimes, the difficulty in XVA modeling is determining whether the exposure to a counterparty at the time of its default is positive or negative. Calculating the exposure requires an underlying valuation model and this is, therefore, a model-dependent term (Karlsson, Jain, & Oosterlee, 2016). A distinction can be made in the class of valuation models, there are models which use a closed-form expression for the exposure at default time, while other models use a simulation-based approach to determine exposure. In the first class, one can derive closed-form expressions by e.g. assuming that the value of a contract is normally distributed (Pykhtin & Rosen, 2010), or using option valuation formulae like the Black formula of Black (1976) to approximate the future value of a contract. Although these approaches avoid constructing a complicated simulation framework, the downside of these approaches is that they do not take netting effects and/or collateral agreements into account (Pykhtin & Rosen, 2010), which are often observed in practice (Joshi & Kwon, 2016; Karlsson, Jain, & Oosterlee, 2016; Zhu &

Lomibao, 2005; Zhu & Pykhtin, 2007).

Alternatively, one could calculate the exposure at the time of default by using a Monte Carlo (MC) simulation method. For example, Brigo et al. (2009) simulate from a CIR interest rate model and combine these simulated interest rates with a CVA calculation for commodity swap contracts. Others use an analogy with pricing European, American, or Bermudan options for the EPE/ENE estimation. These options can be priced using an MC simulation (see e.g. Boyle (1977), Broadie et al. (1997), Joy et al. (1996), and Rogers (2002)) and also require an intermediate valuation, but then to determine whether the continuation value of an option is positive or negative. The analogy can be combined with an iterative Least-Squares regression in the Monte Carlo (LSMC) simulation to estimate the future option value. This methodology originates in Carriere (1996) and is extended in Longstaff and Schwartz (2001) and Karlsson, Jain, and Oosterlee (2016). In short, the idea behind the LSMC simulation method is to approximate the continuation value of an American or Bermudan option by estimating it using backward induction and a set of regression functions in each backward step. However, in a portfolio setting with many derivative trades and intermediate cash flows, this is a time-consuming process.

To tackle the problem of long computation times, improvements are made by e.g. Gamba (2003) and Areal et al. (2008), the latter of which claim to reduce the computation time compared to the original LSMC algorithm by one-third. They test their algorithm on a portfolio of options and show that they can increase the accuracy of the LSMC algorithm while decreasing the computation time. They achieve this by applying several variance reduction techniques and using a different set of regression functions. Applying the LSMC algorithm to a portfolio of options is particularly relevant for this paper, as we also evaluate an entire portfolio. However, these papers use their algorithms to make capital budgeting decisions, which is different from our goal of calculating XVA.

Joshi and Kwon (2016) apply the Longstaff-Schwartz methodology in an XVA setting for Bermudan and cancellable swaps. Their simulation method can be extended to a derivatives portfolio setting that takes netting and collateral agreements into account. Joshi and Kwon (2016) note that the regression functions in the original LSMC algorithm struggle to correctly estimate the future exposure, and they therefore reformulate the CVA formula such that the regression functions only need to be used to determine the sign (positive or negative) of the exposure. The CVAs computed with their algorithm show significantly lower errors than other LSMC algorithms, therefore we implement their XVA methods in this research. However, we use different underlying variables and evaluate a different portfolio of derivative options (as we also allow for carbon forwards in a portfolio).

3 Commodity Interest Rate Market Model

The goal of this research is to calculate Valuation Adjustments (XVA) for carbon forward contracts based on simulated interest rates, exchange rates and convenience rates. To do so, we first need a model that describes their dynamics. The convenience rate of a commodity is the return obtained from holding it (if positive), or the cost of holding it (if negative) (O'Connor et al., 2015). The simulation method is based on the assumption that we can see the carbon emission allowance market as a currency market (Benz & Trück, 2009; Pilz & Schlögl, 2013). This way, we can apply the multicurrency framework of Schlögl (2002) in which two or more currency markets are connected via their exchange rates. This linkage assumption is not arbitrary, as the Commodity LIBOR Market Model (CLMM) of Pilz and Schlögl (2013) uses the same supposition. In this paper, we will assume that the euro is the domestic currency and that all other currencies and the carbon market are the foreign currencies.

The euro interest rate curve and the carbon convenience rate curve can be linked through the euro carbon spot exchange rate. The carbon spot exchange rate (or carbon spot price) indicates how much money needs to be paid right now to receive one unit of carbon emission allowance. This spot exchange rate can be readily observed in the carbon market and is thus treated as given. In this research, we are going to use a Commodity Interest Rate Market model (CIRM model) for the interest and convenience rates. This model is an extension of the CLMM as outlined in Pilz and Schlögl (2013). Note that in the implementation of the model, we will use a discrete version of the CIRM model with displaced forward rates, both for the currency forward rates as well as for the convenience rate. However, to make the model easier to understand we first explain the model in a discrete setting without displacement.

3.1 The Forward Rate Markets

Let the tenor structure of the CIRM model be given by $[T_0, \ldots, T_{N+1}]$. The tenor dates of the CIRM model are consistent with the forward rates with maturity $T_0, \ldots, T_N$. The valuation (or value) time is denoted by time T_0 and thus corresponds to the first tenor for which we simulate rates with the CIRM model. Additionally, we have running time t, which can also correspond to any time in the tenor structure and this will be used in the upcoming section to derive general expressions (e.g. stochastic differential equations and other statements that hold for any tenor in the tenor structure).

Let M denote the number of currencies in the model. Furthermore, let $m = 1$ correspond to the euro, $m = 2, ..., M-1$ to the other monetary currencies, and $m = M$ to the carbon 'currency'. Additionally, we introduce the concept of state variables. In the CIRM model, we have d state variables, which correspond to all interest and exchange rates in the CIRM model. Since we have M currencies and $N+1$ forward rates per currency, we know that the number of state variables equals $d = M \times (N+1) + M - 1 = M \times (N+2) - 1$.

We define $\delta_{T_i}^{(m)}$ to be the year factor for the period $[T_i, T_{i+1}]$ calculated using the appropriate daycount convention for the market of currency m. Additionally, we define $n(t)$ for $t \in [T_0, T_N]$ to be the index such that $T_{n(t)-1} \leq t < T_{n(t)}$. Thus, $n(t)$ is the index of the tenor date that is closest to expiry and therefore can be seen as the spot index. Furthermore, we assume a given probability space $(\Omega, \{\mathcal{F}_t\}_{t\in[T_0,T_{N+1}]}, \mathbb{P}^{T_i})$. Here, we assume a sample space Ω, a filtration $\mathcal{F}_t$ and a T_i-forward probability measure $\mathbb{P}^{T_i}$. We will use this probability space to indicate which of the variables are stochastic or deterministic.

Next, we define the value of a carbon forward contract. To this end, we use that $V_{forward}(t, T_i)$

is the present value at time t of carbon forward contract maturing at time T_i. Let K denote the forward (strike) price which is agreed upon in the carbon forward contract. The carbon forward exchange rate curve, which shows the price for one carbon emission allowance at various points in the future, is an important element in the CIRM model and is discussed in more detail in Section 3.2. Denote the carbon forward exchange rate at time T_i and observed at time t by $X^{(1:M)}(t, T_i)$. Then, we can compute the value of the carbon forward contract in the CIRM model by postulating that the value at time t is given by

$$V_{forward}(t, T_i) = B^{(1)}(t, T_i)\,(X^{(1:M)}(t, T_i) - K), \quad \forall i = 0, \dots, N+1. \tag{3.1}$$

Furthermore, let $B^{(m)}(t, T_i)$ be the value of an arbitrage-free currency m zero-coupon bond at time $t \leq T_i$ that pays 1 at maturity T_i in currency m. Note that $B^{(m)}(T_i, T_i) = 1$ for all m. The carbon bond price $B^{(M)}(t, T_i)$ is the amount of carbon emission allowances that have to be invested at time t to physically receive one unit of carbon emission allowance at time T_i. Note that carbon bonds are not traded on the market, but we will show in Section 3.4 that we do not need carbon bond prices to calibrate the model. Then, the carbon forward rates can be seen as the yields on this carbon bond. The carbon forward rates are sometimes called convenience forward rates in the literature, hence we use these terms interchangeably.

Then, based on zero-coupon bonds, we can compute (convenience) forward rates by using the following relationship between forward rates and bond prices (Schlögl, 2002)

$$F_t^{(m)}(T_i) = \frac{1}{\delta_i}\frac{B^{(m)}(t, T_i)}{B^{(m)}(t, T_{i+1})} - 1, \quad \forall m = 1, \dots, M, \quad \forall i = 0, \dots, N. \tag{3.2}$$

Using this methodology, we obtain forward rates for each tenor date. Define $F_t^{(m)}(T_i)$ to be the time $t \leq T_i$ value of the forward rate corresponding to currency m and spanning the period $[T_i, T_{i+1}]$. Therefore, when speaking of forward rates, we also mean the carbon (forward) convenience rates from here onwards. To describe the dynamics of the forward rates, we need to make an assumption on their distribution. One of the possibilities is outlined in Schlögl (2002). In his first scenario, Schlögl (2002) assumes that all forward rates are lognormally distributed under their respective $\mathbb{Q}_{(m),T_{i+1}}$ forward measure. Observe that $\mathbb{Q}_{(1),T_{i+1}}$ is called the domestic forward measure, since it corresponds to the euro market. The dynamics of all forward rates are given by

$$dF_t^{(m)}(T_i) = F_t^{(m)}(T_i)\,\sigma_t^{(m)}(T_i) \cdot dW_t^{\mathbb{Q}_{(m),T_{i+1}}}, \quad t \in [0, T_i], \quad \forall i = 0, \dots, N, \tag{3.3}$$

where $\sigma_t^{(m)}(T_i) : [0, T_i] \to \mathbb{R}_{\geq 0}$ and $F_t^{(m)}(T_i) : \Omega \times [0, T_i] \to \mathbb{R}$. In Equation (3.3), $\sigma_t^{(m)}(T_i)$ denotes the at time t observed d-dimensional instantaneous volatility vector of the currency m forward rates covering the period $[T_i, T_{i+1}]$. Note that $\sigma_t^{(m)}(T_i)$ is a deterministic parameter, while the forward rate is a stochastic parameter due to the Brownian motion in Equation 3.3 so the fact that these parameters have the same subscript does not imply anything about their randomness. Furthermore, we are not the first to add the observation time to the volatility parameters, as e.g. Karlsson, Pilz, and Schlögl (2016) and Pilz and Schlögl (2013) do the same for their (potentially) deterministic forward rate volatilities. Our reason for explicitly adding the observation time as a subscript is that we need observation time-dependent volatilities for the CIRM model calibration as described in Section 4. From here on, we will not explicitly state anymore that the (stochastic) differential equations hold for all $t \in [0, T_i]$ and $i = 0, \ldots, N$, but we remark that this holds unless stated otherwise.

Furthermore, $W_t^{\mathbb{Q}_{(m),T_{i+1}}}$ is a d-dimensional standard Brownian motion with independent elements under the $\mathbb{Q}_{(m),T_{i+1}}$ probability measure. Thus, we have $W_t^{\mathbb{Q}_{(m),T_{i+1}}} = (W_{1,t}^{\mathbb{Q}_{(m),T_{i+1}}}, ..., W_{d,t}^{\mathbb{Q}_{(m),T_{i+1}}})$ where $W_{j,t}^{\mathbb{Q}_{(m),T_{i+1}}}$ for each $1 \leq j \leq d$ is a 1-dimensional Brownian motion. These 1-dimensional Brownian motions are correlated to one another, and we denote the correlation between the two 1-dimensional Brownian motions $W_{e,t}^{\mathbb{Q}_{(m),T_{i+1}}}$ and $W_{f,t}^{\mathbb{Q}_{(m),T_{i+1}}}$ by $\rho_{e,f,t}$ for $1 \leq e, f \leq d, e \neq f$. To calibrate these correlation parameters we use the methodology outlined in Section 4.3.

3.1.1 Displaced Forward Rates

Due to the log-normality assumption that is used in both the CLMM of Pilz and Schlögl (2013) as well as the multicurrency model of Schlögl (2002), it is not possible to have negative interest rates. However, recent behavior of interest rates shows that this phenomenon can actually occur, hence a forward interest rate model should allow for this. To do so, we denote the displaced simple forward rates by

$$\dot{F}_t^{(m)}(T_i) = F_t^{(m)}(T_i) + \Lambda_i^{(m)}, \tag{3.4}$$

with $\Lambda_i^{(m)} \in \mathbb{R}$ being a deterministic and time-invariant parameter for each period $[T_i, T_{i+1}]$ that ensures that the initial displaced forward rates for currency m are positive. This way, the displaced forward rates are assumed to be log-normally distributed as well. In this research, we simulate $\dot{F}_t^{(m)}(T_i)$ from a log-normal distribution and then subtract a value of $\Lambda_i^{(m)}$ to obtain forward rates $F_t^{(m)}(T_i)$ that can be negative. Note that we also perform this displacement methodology for the convenience rates, as Pilz and Schlögl (2013) note that convenience rates can also be negative.

3.2 The Forward Exchange Rate Markets

One of the key aspects of the CIRM model is the forward exchange rate $X^{(1:m)}(t, T_{i+1})$. This is the forward exchange rate at time T_{i+1} that denotes the time t value of one unit of currency m expressed in units of euros, without loss of generality. Note that $X^{(1:M)}(t, T_{i+1})$ is the carbon forward price, as this is the price in euros of 1 carbon emission allowance at time T_{i+1}. The exchange rate can be used as a so-called measure link that connects the euro forward interest rate curve to other forward curves. An illustrative picture of the measure link concept is shown in Figure 1.

Figure 1. Relationships Between Different Tenors in the CIRM Model

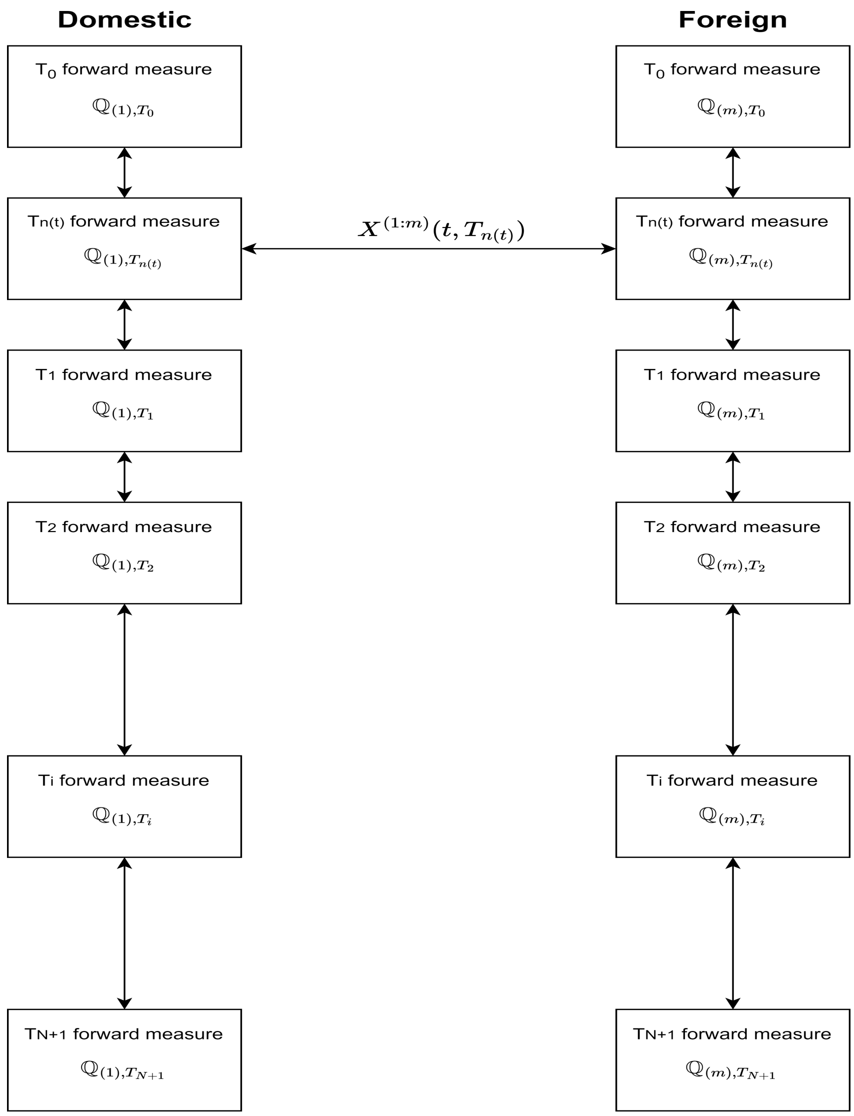

Note. Figure shows the setup of the tenor structure in the CIRM model. For illustrative purposes, we only depict the relationships in case we have two currencies ('domestic' and 'foreign') but this can be extended to a situation with multiple foreign currencies. In our case, we see the euro as the domestic currency and the foreign currency can be equal to any currency m. The tenor structures of different currencies are linked through the spot forward exchange rate as shown by the horizontal arrow.

Thus, by using $X^{(1:M)}(t, T_{i+1})$ as a link between the euro forward market and the convenience forward market, we can simulate a convenience forward curve. The derivation of the convenience forward curve is outlined in Section 3.4.

To compute the forward exchange rates, we use the following relationship as described in Pilz and Schlögl (2013)

$$X^{(1:m)}(t, T_{i+1}) = \frac{X^{(1:m)}(t) \cdot B^{(m)}(t, T_{i+1})}{B^{(1)}(t, T_{i+1})}, \tag{3.5}$$

for $i = 0, \ldots, N$ and $m = 2, \ldots, M$. $X^{(1:m)}(t)$ denotes the spot exchange rate from currency m to euros. Here, we used the definitions of the zero coupon bonds in the carbon/foreign currency market, $B^{(m)}(t, T_{i+1})$, and the domestic currency market, $B^{(1)}(t, T_{i+1})$, as given in Section 3.1. However, we choose to simulate the spot forward exchange rate $X^{(1:m)}(t, T_{n(t)})$ instead of the spot exchange rate. For this reason, we also use the following definition of the forward exchange rate in this paper

$$X^{(1:m)}(t, T_{i+1}) = X^{(1:m)}(t, T_{n(t)}) \frac{B^{(m)}(t, T_{n(t)}, T_{i+1})}{B^{(1)}(t, T_{n(t)}, T_{i+1})}, \tag{3.6}$$

where the spot forward exchange rate $X^{(1:m)}(t, T_{n(t)})$ is defined as in Equation (3.5) and $B^{(m)}(t, T_{n(t)}, T_{i+1}) = \frac{B^{(m)}(t,T_{i+1})}{B^{(m)}(t,T_{n(t)})}$ is the time t value of a bond in currency m which starts at $T_{n(t)}$ and ends at T_{i+1}. It is also convenient to move from one forward exchange rate to another in the CIRM model. Equation (3.6) allows us to do so by stating that any forward exchange rate is a function of the spot forward exchange rate and two zero coupon bonds.

Looking closer at Equation (3.6), we see that the numerator in the equation is defined as a foreign bond price denoted in the domestic currency at the spot forward exchange rate. Hence, this makes it a domestic and tradeable asset. No-arbitrage rules imply that a tradeable asset divided by a numéraire ($B^{(1)}(t, T_{n(t)}, T_{i+1})$ in this case) should be a (lognormal) martingale under the corresponding forward measure $\mathbb{Q}_{(1),T_{i+1}}$ (Schlögl, 2002). For an explanation of the numéraire concept see Appendix A. The assumption that the forward exchange rate is a lognormal martingale implies that the conditional expectation of the exchange rate at any point in time is equal to its present value, irrespective of its previous values. Because the conditional expectation of the forward exchange rate is equal to its present value, there is no possibility to obtain arbitrage opportunities. The lognormality assumption indicates that the risk-neutral probabilities that follow from $\mathbb{Q}_{(1),T_{i+1}}$ are lognormally distributed. Hence, assuming no arbitrage is equivalent to assuming that the forward exchange rate is a lognormal martingale under a risk-neutral probability measure. We allow ourselves to make this assumption as it is in line with the (first) fundamental theorem of asset pricing. Due to their lognormality, the dynamics of the forward exchange rates under the $\mathbb{Q}_{(1),T_{i+1}}$ probability measure are given by

$$dX^{(1:m)}(t, T_{i+1}) = X^{(1:m)}(t, T_{i+1}) \sigma_t^{(1:m)}(T_{i+1}) \, dW_t^{\mathbb{Q}_{(1),T_{i+1}}}, \tag{3.7}$$

where $\sigma_t^{(1:m)}(T_{i+1}) : [0, T_{i+1}] \to \mathbb{R}_{\geq 0}$ and $X^{(1:m)}(t, T_{i+1}) : \Omega \times [0, T_{i+1}] \to \mathbb{R}$. We denote the volatility of the T_{i+1} forward exchange rate as $\sigma_t^{(1:m)}(T_{i+1})$. Note that due to this setup with lognormal forward rates under their respective forward measures, the forward exchange rates with maturity T_{i+1}, $X^{(1:m)}(t, T_{i+1})$, are only lognormally distributed under the $\mathbb{Q}_{(1),T_{i+1}}$ measure. Lastly, we emphasize that the forward exchange rates are stochastic due to the Brownian motion in Equation (3.7), while the volatility $\sigma_t^{(1:m)}(T_{i+1})$ is a deterministic parameter observed at time t. This is mathematically shown in their definitions in the probability space.

3.3 Forward and Exchange Rates under the Domestic Spot Measure

To simulate from the CIRM model, we need all state variables expressed under the same domestic measure. Fortunately, we have already derived the dynamics of the domestic forward interest and forward exchange rates under the $\mathbb{Q}_{(1),T_{i+1}}$ measure. To achieve the same for the foreign forward rates, we use that the equivalent foreign forward measures and domestic forward measures are linked by the Radon-Nikodym derivative as

$$\frac{dW_t^{\mathbb{Q}_{(m),T_{i+1}}}}{dW_t^{\mathbb{Q}_{(1),T_{i+1}}}} = \exp\left(\int_0^{T_{i+1}} \sigma_t^{(m)}(T_i)\, dW_t^{\mathbb{Q}_{(1),T_{i+1}}} - \frac{1}{2}\int_0^{T_{i+1}} \sigma_t^{(m)}(T_i)^2\, dt\right). \tag{3.8}$$

Then, by applying the Cameron-Martin-Girsanov theorem we find that the foreign forward measure can be expressed in terms of the domestic forward measure as

$$dW_t^{\mathbb{Q}_{(m),T_{i+1}}} = dW_t^{\mathbb{Q}_{(1),T_{i+1}}} - \sigma_t^{(1:m)}(T_{i+1})\, dt. \tag{3.9}$$

Using this result, we obtain the following expression for the dynamics of the foreign forward interest rates

$$\begin{aligned} dF_t^{(m)}(T_i) &= \dot{F}_t^{(m)}(T_i)\sigma_t^{(m)}(T_i) \cdot dW_t^{\mathbb{Q}_{(m),T_{i+1}}} \\ &= \dot{F}_t^{(m)}(T_i)\left(-\sigma_t^{(m)}(T_i)\cdot\sigma_t^{(1:m)}(T_{i+1})\, dt + \sigma_t^{(m)}(T_i)\, dW_t^{\mathbb{Q}_{(1),T_{i+1}}}\right), \end{aligned} \tag{3.10}$$

for all currencies $m = 2, \ldots, M$. Note that the domestic (i.e. euro) forward rates are already described under the domestic forward measure and thus do not require a change of measure.

Lastly, the only exchange rates that we simulate using the CIRM model are the spot forward exchange rates $X^{(1:m)}(t, T_{n(t)})$. This is because all other forward exchange rates can be computed using these rates as can be seen in Equation (3.6). To simulate the spot forward exchange rates, we need the dynamics of all rates to be given under the same domestic spot measure. To shorten notation, we define the domestic spot measure to be $\mathbb{Q}_{(1)} = \mathbb{Q}_{(1),n(t)}$. To convert rates under the $\mathbb{Q}_{(1),T_{i+1}}$ measure to rates under the $\mathbb{Q}_{(1)}$ measure, we use the following relationship as given by

$$dW_t^{\mathbb{Q}_{(1),T_{i+1}}} = dW_t^{\mathbb{Q}_{(1)}} + \sum_{k=n(t)}^{i} \frac{\delta_{T_k}^{(1)}(\dot{F}_t^{(1)}(T_k))}{1+\delta_{T_k}^{(1)}\dot{F}_t^{(1)}(T_k)}\sigma_t^{(1)}(T_k)\, dt, \tag{3.11}$$

which is similar to the drift term as given in Equation (4) in Schlögl (2002). This relationship holds both for forward rates as well as for forward exchange rates. Plugging in Equation (3.11) in Equation (3.10) we get

$$\begin{aligned} dF_t^{(m)}(T_i) = \dot{F}_t^{(m)}(T_i)\Big(\big(&-\sigma_t^{(m)}(T_i)\cdot\sigma_t^{(1:m)}(T_{i+1}) \\ &+ \sigma_t^{(m)}(T_i)\sum_{k=n(t)}^{i}\frac{\delta_{T_k}^{(1)}(\dot{F}_t^{(1)}(T_k))}{1+\delta_{T_k}^{(1)}\dot{F}_t^{(1)}(T_k)}\sigma_t^{(1)}(T_k)\big)dt + \sigma_t^{(m)}(T_i)\, dW_t^{\mathbb{Q}_{(1)}}\Big), \end{aligned} \tag{3.12}$$

for the dynamics of all forward rates under the domestic spot measure.

For the spot forward exchange rates, it suffices to give the dynamics of the spot forward exchange rates as a specific version of Equation (3.7) and using $\mathbb{Q}_{(1)} = \mathbb{Q}_{(1),n(t)}$, as we then get

$$dX^{(1:m)}(t, T_{n(t)}) = X^{(1:m)}(t, T_{n(t)})\left(\sigma_t^{(1:m)}(T_{n(t)})\, dW_t^{\mathbb{Q}_{(1)}}\right). \tag{3.13}$$

3.4 Constructing the Convenience Forward Rates

We use the CIRM model to simulate, among other things, convenience forward rates $F_t^{(M)}(T_i)$. We simulate forward rates for each tenor in the tenor structure. To do so, we first compute a convenience discount factor for each tenor. Then, we calculate the convenience forward rates based on these discount factors. Note that we will explain the methodology for maturity T_i but it can be analogously applied for other maturities in $[T_0, \ldots, T_{N+1}]$.

We start by determining the convenience discount factor with maturity T_i. For this, we use that the value of a forward contract with maturity T_i can be calculated in two ways. One way to compute it is by postulating that the value of a forward contract at time t is given by Equation (3.1). Alternatively, we can express the value of the carbon forward contract as a function of the carbon spot exchange rate. To this end, we use the definition of $X^{(1:M)}(t, T_i)$ as shown in Equation (3.5) to get

$$\begin{aligned} V_{forward}(t, T_i) &= B^{(1)}(t, T_i) \left(\frac{X^{(1:M)}(t) \cdot B^{(M)}(t, T_i)}{B^{(1)}(t, T_i)} - K \right), \\ &= X^{(1:M)}(t) \cdot B^{(M)}(t, T_i) - B^{(1)}(t, T_i) K. \end{aligned} \tag{3.14}$$

By construction, the value of the carbon forward contract in Equation (3.1) should be the same as the value implied by Equation (3.14). Equating these two formulations gives us an expression for the convenience discount factor for tenor T_i as

$$\begin{aligned} X^{(1:M)}(t) \cdot B^{(M)}(t, T_i) - B^{(1)}(t, T_i) K &= B^{(1)}(t, T_i) (X^{(1:M)}(t, T_i) - K), \\ \Leftrightarrow X^{(1:M)}(t) \cdot B^{(M)}(t, T_i) &= B^{(1)}(t, T_i) X^{(1:M)}(t, T_i), \\ \Leftrightarrow B^{(M)}(t, T_i) &= \frac{B^{(1)}(t, T_i) X^{(1:M)}(t, T_i)}{X^{(1:M)}(t)}, \end{aligned} \tag{3.15}$$

where $B^{(M)}(t, T_i)$ denotes the convenience discount factor that matures at time T_i. As mentioned before, we can compute the convenience discount factor for any tenor in the tenor structure. Based on the convenience discount factors, we calculate the convenience forward rates using Equation (3.2). To sum up, in this section we have derived the dynamics under the euro spot measure of all rates in the CIRM model. Also, we have described how we are going to compute the convenience forward rates. Next, to simulate forward interest and forward exchange rates we still need to specify their instantaneous volatilities, which we will discuss in the next section.

4 Calibration of the Interest Rate Market Model

To use the Commodity Interest Rate Market model (CIRM model) in the context of carbon forward contracts, we first need to calibrate the model. We start this section by describing the calibration of the instantaneous volatilities of the forward rates ($\sigma_t^{(m)}(T_i)$) in Section 4.1. This calibration is done based on market data on interest rate caps and floors. After that, we discuss the calibration of the forward exchange rate volatilities ($\sigma_t^{(1:m)}(T_i)$) and of the correlations between all forward interest and forward exchange rate volatilities. We discuss the exchange rate volatility calibration in Section 4.2 and the correlation calibration in Section 4.3.

4.1 Calibration of Interest Rate Volatilities

To calibrate the instantaneous volatilities of the forward interest rates, we use market data on prices of caps and floors on the interest rate of currency m. A cap (floor) is a financial product that provides protection against rising (falling) interest rates to its buyer. It provides payment at the end of each of a pre-determined set of periods when the floating interest rate exceeds the cap rate (or below the floor rate in case of a floor). This way, the buyer can limit its losses in case of large movements in the underlying interest rate. From here on, we specifically discuss caps, but floors can be used analogously for the interest rate volatility calibration.

A cap can be seen as a series of European call options on the interest rate which all have different maturity dates. Each individual call option is also referred to as a caplet in this context. At the maturity of each caplet, the buyer of the cap receives a payment in case the floating interest rate exceeds the cap rate, thus ensuring that the buyer of the cap never pays an interest rate above the cap rate. The payoff of a caplet on currency m's interest rate with maturity date T_{i+1} can be described as follows

$$PF^{(m)}_{caplet,T_{i+1}} = \delta^{(m)}_{T_i}(F^{(m)}_{T_i}(T_i) - K^{(m)})^+, \tag{4.1}$$

where $K^{(m)}$ denotes the cap rate of the cap contract. Similar as for the CIRM model, we again assume that forward rates are lognormally distributed. Hence, we can use the Black 76 formula (Black, 1976) to price individual caplets. The price of a caplet on currency m's interest rate implied by Black's formula at time t is given as

$$\begin{aligned} V^{(m)}_{caplet,t}(T_i, \dot{K}^{(m)}, \sigma^{(m)}_t(T_i)) &= \delta_{T_i} B^{(m)}(t, T_{i+1}) \left(\dot{F}^{(m)}_t(T_i)\Phi(d_1) - \dot{K}^{(m)}\Phi(d_2)\right), \\ d_1 &= \frac{\ln\left(\frac{\dot{F}^{(m)}_t(T_i)}{\dot{K}^{(m)}}\right) + \frac{1}{2}\sigma^{(m)}_t(T_i)^2 T_i}{\sigma^{(m)}_t(T_i)\sqrt{T_i}}, \\ d_2 &= d_1 - \sigma^{(m)}_t(T_i)\sqrt{T_i}, \end{aligned} \tag{4.2}$$

where $\dot{K}^{(m)}$ denotes the displaced cap rate as the forward rate is also displaced and $\Phi(\cdot)$ is the standard normal distribution. As can be seen, the price of a caplet depends on its cap rate, time-to-maturity, and the implied volatility of the forward rate. Out of these parameters, only the implied volatility is unknown. We can observe prices on interest rate caps in the market. From these cap prices, we can derive caplet prices by performing a global caplet stripping procedure. Using the caplet prices and inverting Black's formula, we can calculate the implied volatilities using Equation (4.2). Note that by using Black's formula for the implied volatilities, we assume that the volatility is constant over the lifetime of the cap. This means that, for example, the volatility implied by the cap extending over the period $[0, T_2]$ has the same volatility for the period $[0, T_1]$ as for the period $[T_1, T_2]$. Although this is not a very realistic assumption, this paper is not the first to make it. Brigo and Mercurio (2006) also consider constant forward interest rate volatilities for their calibration of a LIBOR Market Model.

An advantage of using Black implied volatilities is that they are independent of the cap rate of the cap, thus any cap on a currency can, in theory, be used. Another benefit is that there exists an analytical solution for the Black implied volatility, which simplifies the calibration procedure compared to using stochastic volatility models with numerical solutions. Lastly, note that we use at-the-money caps to calibrate the volatilities, and we do not allow for a volatility smile due to the assumption that the volatility is independent of the cap rate. This can be seen as a limitation of this research, but we believe that for the purpose of this paper, namely simulating valuation adjustments, the simplicity of this calibration method outweighs the potential benefits of using more complicated calibration procedures.

To generate from the CIRM model, we take the volatilities of the carbon forward interest rates (which together make the convenience curve) equal to zero. Although this might sound counter-intuitive, specifying this parameter equal to zero does not lead to unwanted behavior. By doing this, we obtain convenience forward rates that are deterministic and constant, as follows from Equation (3.12) for the dynamics of the forward rates. Note that they are deterministic and constant in the sense that we construct the convenience forward curve only once (at T_0) using Equation (3.15), and after that the shape of the convenience curve does not change anymore. However, this does not mean that the convenience forward rates are the same for each period, e.g. the convenience forward rate for the period $[T_0, T_1]$ need not be equal to the convenience forward rate between $[T_1, T_2]$. This follows from Equation (3.15), where we see that the convenience discount factor depends on the euro forward interest discount factor and the carbon spot and forward exchange rate, which do not have zero volatility and thus vary over time.

We assume zero carbon forward interest rate volatility because the aforementioned calibration method for interest rate volatilities would require using data on carbon caps and floors, which do not exist. Moreover, assuming a specific level for carbon interest rate volatility may prove difficult, as carbon bonds also do not exist and there are no closely related asset classes that do have bonds available. Therefore, we assume that the carbon interest rate volatility is equal to zero. This does not lead to miscalibration of the CIRM model, as all carbon-related volatility in the market is captured in the volatility of the carbon exchange rate, which we do imply from the market. The additional benefit of assuming a level of zero volatility is that several correlation parameters related to carbon also become equal to zero, which makes for an easier estimation of the CIRM model parameters.

4.2 Calibrating the Forward Exchange Rate Volatility

We calibrate the forward exchange rate volatilities using a different procedure. We explain the procedure under the assumption that all forward interest rate volatilities are equal to zero. We only make this assumption in this section, such that we can simplify the explanation of the forward exchange rate volatility procedure. In practice, forward interest rate volatilities are not equal to zero, and non-zero interest rate volatilities lead to an additional component in the exchange rate volatility. As the focus of this paper is to calculate XVA and not the calibration of the CIRM model, we believe that an extensive derivation of the exchange rate volatility as a function of interest rate volatilities is out of the scope of this paper.

We start the explanation by referring back to Equation (3.6), where we defined the forward exchange rate as a function of the spot forward exchange rate and two discount bonds. Based on this equation, we see that the forward exchange rate is related to the spot forward exchange rate using a twice continuously differentiable function. We will later see that this characteristic is convenient because we will apply Itô's lemma to the spot forward exchange rate.

As noted in Section 3.2, the spot forward exchange rate is a martingale. This means that we can write its dynamics under the domestic spot measure as given in Equation (3.13), while the dynamics of the forward exchange rate $X^{(1:m)}(t, T_i)$ under the same measure can be written as

$$dX^{(1:m)}(t, T_{i+1}) = X^{(1:m)}(t, T_{i+1}) \left(\ldots dt + \sigma_t^{(1:m)}(T_{i+1}) \, dW_t^{\mathbb{Q}_{(1)}} \right), \tag{4.3}$$

where we intentionally left out the drift term as we are looking for an expression for the volatility parameter, which is determined by the stochastic component of the differential equation only. Next, we apply Itô's lemma to derive the dynamics of the spot forward exchange rate under the domestic spot measure and use the assumption that interest rate volatilities are zero to get

$$dX^{(1:m)}(t, T_{i+1}) = X^{(1:m)}(t, T_{i+1}) \left(\ldots dt + \frac{\partial X^{(1:m)}(t, T_{i+1})}{\partial X^{(1:m)}\left(t, T_{n(t)}\right)} \frac{dX^{(1:m)}\left(t, T_{n(t)}\right)}{X^{(1:m)}(t, T_{i+1})} \right). \tag{4.4}$$

We can write the term in brackets that is not related to the drift as

$$\begin{aligned}
&\frac{\partial X^{(1:m)}(t, T_{i+1})}{\partial X^{(1:m)}\left(t, T_{n(t)}\right)} \frac{dX^{(1:m)}\left(t, T_{n(t)}\right)}{X^{(1:m)}(t, T_{i+1})} = \frac{B^{(m)}(t, T_{n(t)}, T_{i+1})}{B^{(1)}(t, T_{n(t)}, T_{i+1})} \cdot \\
&\frac{X^{(1:m)}\left(t, T_{n(t)}\right) \left(\sigma_t^{(1:m)}\left(T_{n(t)}\right) dW_t^{\mathbb{Q}_{(1)}}\right)}{X^{(1:m)}(t, T_{i+1})}, \\
&= \frac{B^{(1)}(t, T_{n(t)}, T_{i+1}) X^{(1:m)}(t, T_{i+1})}{B^{(1)}(t, T_{n(t)}, T_{i+1}) X^{(1:m)}(t, T_{i+1})} \sigma_t^{(1:m)}\left(T_{n(t)}\right) dW_t^{\mathbb{Q}_{(1)}}, \\
&= \sigma_t^{(1:m)}\left(T_{n(t)}\right) dW_t^{\mathbb{Q}_{(1)}},
\end{aligned} \tag{4.5}$$

where we utilized the expressions for the forward exchange rate, the dynamics of the spot forward exchange rate, and a rewritten version of Equation (3.6) to obtain this result. Then, using this expression in Equation (4.4) we obtain

$$dX^{(1:m)}(t, T_{i+1}) = X^{(1:m)}(t, T_{i+1}) \left(\ldots dt + \sigma_t^{(1:m)}\left(T_{n(t)}\right) dW_t^{\mathbb{Q}_{(1)}} \right). \tag{4.6}$$

Comparing this to the expression for the forward exchange rate dynamics in Equation (4.3), we can extract the relationship between the volatility of the forward exchange rate and the volatility of the spot forward exchange rate. Next, we assume that the forward exchange rate volatility is piecewise constant. Hence, we have constant volatility (and variance) over the period $[T_i, T_{i+1})$. The forward exchange rate volatility over the period $[T_i, T_{i+1})$ is then given by

$$\sigma_t^{(1:m)}(T_{i+1}) = \sigma_t^{(1:m)}\left(T_{n(t)}\right), \tag{4.7}$$

and this holds for every $i = 0, \ldots, N$. It is important to note that this does not imply that the volatility is always at the same constant level. Specifically, the forward exchange rate volatility depends on observation time t too, hence the volatility at one observation time need not be the same as the volatility for another observation time because the volatility is only piecewise constant. Therefore, it is essential for the volatility parameters to have a subscript indicating the observation time. Given the result in Equation (4.7), it is sufficient to derive an expression for $\sigma_t^{(1:m)}\left(T_{n(t)}\right)$ to calibrate the forward exchange rate volatilities of the CIRM model. Following our previous assumption, we use that also the spot forward exchange rate volatility is piecewise constant to get

$$\begin{aligned} \sigma_t^{(1:m)}\left(T_{n(t)}\right) :=& \sigma_{T_0}^{(1:m)}(T_1)\, \mathbb{1}_{T_0 \leq t < T_1} + \cdots + \sigma_{T_N}^{(1:m)}(T_{N+1})\, \mathbb{1}_{T_N \leq t < T_{N+1}}, \\ =& \sum_{i=0}^{N} \sigma_{T_i}^{(1:m)}(T_{i+1})\, \mathbb{1}_{T_i \leq t < T_{i+1}}, \end{aligned} \tag{4.8}$$

which follows from the fact that the value of the spot forward exchange rate volatility depends on when the spot forward volatility is observed (i.e. between which tenors), and from the result in Equation (4.7).

Then, we are left with the calibration of the individual volatility parameters of Equation (4.8). To calibrate these, we make use of the at valuation time T_0 observed exchange rate volatility implied by the market ($\sigma_{T_0,mkt}^{(1:m)}(T_{i+1})$). So the valuation time corresponds to the first tenor for which we have simulated state variables. This volatility implied by the market spans the observational period $[T_0, T_{i+1}]$ for the volatility relating to the T_{i+1} tenor date and stays at a constant level during this observational period. By contrast, $\sigma_t^{(1:m)}(T_{i+1})$ describes an observation time-varying forward exchange rate volatility for the T_{i+1} tenor date. To establish the relationship between the market volatility and the volatility used in the CIRM model, we follow the same approach as Pilz and Schlögl (2013) and state that the integral over the time-dependent instantaneous squared volatilities over the period $[T_0, T_{i+1}]$ should be equal to the market-implied variance, i.e.

$$\begin{aligned} (\sigma_{T_0,mkt}^{(1:m)}(T_{i+1}))^2 T_{i+1} &= \int_{T_0}^{T_{i+1}} (\sigma_s^{(1:m)}(T_{i+1}))^2 ds, \\ &= \int_{T_0}^{T_{i+1}} (\sigma_s^{(1:m)}\left(T_{n(s)}\right))^2 ds, \\ &= \sum_{y=0}^{i} \left(\int_{T_y}^{T_{y+1}} (\sigma_s^{(1:m)}\left(T_{n(s)}\right))^2 ds \right), \\ &= \sum_{y=0}^{i} \left((\sigma_{T_y}^{(1:m)}\left(T_{n(T_y)}\right))^2 (T_{y+1} - T_y) \right), \end{aligned} \tag{4.9}$$

where we used Equation (4.7) and that the volatility is piecewise constant between two subsequent tenors and thus the volatility can be observed at any point in time between two tenors.

We choose to observe the volatility parameter at the beginning of the period between two tenors, as indicated by the subscript of the volatility parameter. When we split the last entry from the sum and again use Equation (4.7), we find the following expression for the forward exchange rate volatility

$$\sigma_{T_i}^{(1:m)}(T_{i+1}) = \sqrt{\frac{1}{T_{i+1}-T_i}\left(\left(\sigma_{T_0,mkt}^{(1:m)}(T_{i+1})\right)^2 T_{i+1} - \sum_{y=0}^{i-1}\left((\sigma_{T_y}^{(1:m)}(T_{n(T_y)}))^2(T_{y+1}-T_y)\right)\right)}, \quad (4.10)$$

so we obtain the at time T_i observed volatility for tenor T_{i+1} using a forward iterative procedure, where we first compute the piecewise constant volatility for the period $[T_0, T_1)$, then the volatility for the period $[T_1, T_2)$ and continue this approach until we can compute the volatility observed at time T_i for the last period $[T_i, T_{i+1})$. As an example, we graphically show how the piecewise constant forward exchange rate volatility is related to the $T_{i+1} = T_5$ tenor date could look like when observing it at different time points in Figure 2.

Figure 2. Piecewise Constant Forward Exchange Rate Volatility

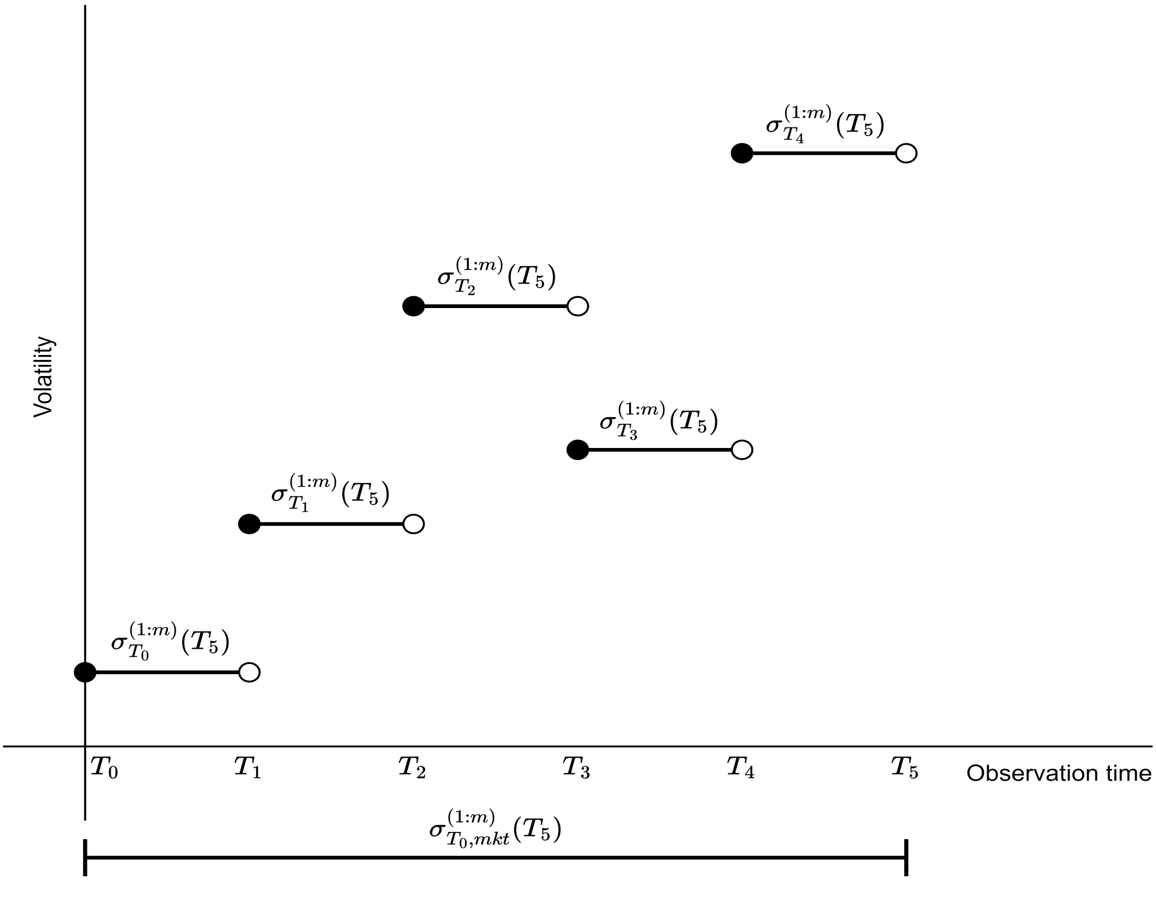

Note. Figure gives an example on how the piecewise constant volatility for a forward exchange rate related to tenor date T_5 could look like. Note that no results are shown in this figure, it merely serves as an illustration. Furthermore, to indicate the relationship between the forward exchange rate volatility and the market implied volatility related to the same tenor date, we show what period is spanned by the market volatility below the graph.

To conclude, we can thus calibrate the deterministic forward exchange rate volatility parameters by performing a forward iterative procedure, and making convenient use of the assumption that the forward exchange rate volatility is piecewise constant between two tenor dates.

4.3 Calibrating the Correlations

Besides volatility parameters, we also need correlation estimates for the calibration of the CIRM model. Their usefulness might not be as obvious as for the volatilities since none of the stochastic differential equations in Section 3 contain correlation parameters, but they are an important factor for the CIRM model. This is because the d underlying Brownian motions in the domestic spot measure $\mathbb{Q}_{(1)}$ are assumed to be correlated. Also, the correlation comes back in the differential equations of Section 3 that contain the inner products of two d-dimensional volatility parameters.

Before explaining the correlation calibration method, we distinguish three different classes of correlations. The first class is the correlation between the forward rates of one currency, $F_0^{(a)}$, and the forward rates of another currency, $F_0^{(l)}$, where $a \neq l$ (hereafter called IR-IR correlation). Note that the subscript 0 indicates the time at which the rate is observed. Second, we have the correlation between forward exchange rates corresponding to two currencies $X^{(a:l)}(0, \cdot)$ and a set of forward exchange rates between two other currencies $X^{(r:p)}(0, \cdot)$ where we cannot have that $a = r$ and $l = p$ at the same time (FX-FX correlation). Lastly, we consider the correlation between forward rates of one currency and the forward exchange rates between two currencies (IR-FX correlation). Since the tenor structure consists of multiple forward rates that can all be observed in the market over time, we need to choose one (forward) rate curve as a representative for each particular FX or IR curve. We base our choice on the liquidity of market quotes as we choose the most liquid rates. Denote the representative curve for the currency a interest rate by $F^{(a)}(T_w)$ and the representative curve for the FX rate between currencies a and l by $X^{(a:l)}(\cdot, T_v)$ with times to maturity T_w and T_v, respectively. For all correlations, we use the Pearson correlation coefficient as introduced by Pearson (1895). The Pearson correlation coefficient is estimated using H historical observations per state variable. The time between successive observations is always one business day. Notice that the dynamics of the CIRM model are all given under the risk-neutral measure, while the previously described methodology for the correlation parameters yields results under the real-world measure. Although this is not ideal, the ease of computation that the real-world methodology offers outweighs the additional accuracy of the risk-neutral methodology. Furthermore, estimating correlation parameters under the risk-neutral measure would require the usage of spread option data, which is not available to us. This paper is not the first to choose real-world over risk-neutral estimation for correlation parameters in a LIBOR Market Model, as the closely related Commodity LIBOR Market Model of Pilz and Schlögl (2013) is calibrated using real-world correlation estimates too.

Following the assumptions of the CIRM model, we have that all IR and FX rates are lognormally distributed. However, for the computation of the Pearson correlation estimates we need time series that are normally distributed. Therefore, we compute the log returns of the FX and IR rates and use these for the correlation estimation, as the log returns are normally distributed and a stationary time series.

Then, define $f^{(m)}$ to be the time series consisting of $H-1$ observations on the first differenced forward rate representative for currency m and $x^{(1:m)}$ the same for the differenced forward exchange rate representative. Then, for each pair of forward rates and/or forward exchange rates, we can compute the correlation coefficient as follows

$$\rho(x^{(1:m)}, f^{(m)}) = \frac{\sum_{j=1}^{H-1} \left(x_j^{(1:m)} - \bar{x}^{(1:m)}\right)\left(f_j^{(m)} - \bar{f}^{(m)}\right)}{\sqrt{\sum_{j=1}^{H-1} \left(x_j^{(1:m)} - \bar{x}^{(1:m)}\right)^2}\sqrt{\sum_{j=1}^{H-1} \left(f_j^{(m)} - \bar{f}^{(m)}\right)^2}}, \tag{4.11}$$

where $f_j^{(m)}$ denotes the j-th observation on the forward rate representative, $x_j^{(1:m)}$ the j-th observation on the exchange rate representative and $\bar{f}^{(m)}$ and $\bar{x}^{(1:m)}$ denote the sample mean of the transformed IR and FX representatives, respectively. Equation (4.11) shows how to compute the IR-FX correlation, but the IR-IR and FX-FX correlation can be calculated analogously.

5 XVA

In this section, we elaborate on the computation of Valuation Adjustments (XVA). In this research, we compute XVA on a portfolio level. Therefore, $V(t)$ denotes the netted value at time t of a possibly collateralized portfolio of derivative contracts which include, among other things, carbon forward contracts. The concepts of netting and collateralization are explained in Section 5.3. First, we dive into the calculation of the Credit Valuation Adjustment (CVA). After that, we discuss the closely related Debt Valuation Adjustment (DVA) computation. After the XVA explanations, we quickly touch upon the concept of netting and collateral agreements in Section 5.3 and, lastly, we dive into the construction of the analytical benchmark model in Section 5.4.

5.1 Credit Valuation Adjustment

As mentioned in Section 1, the CVA is the adjustment to the value of a portfolio to account for counterparty default risk. The CVA formula consists of three elements, the Expected Positive Exposure (EPE), the Loss Given Default (LGD), and the Probability to Default (PD). Two versions of CVA exist, unilateral and bilateral CVA. Unilateral CVA assumes that the institution that computes the CVA cannot default, while bilateral CVA also takes the survival probability of the institution into account (Rosen & Saunders, 2012). Since the assumptions of bilateral CVA are more realistic, we decide to compute the bilateral CVA. Furthermore, we assume that the exposure to a counterparty is independent of the counterparty's PD, so we assume that no wrong-way and right-way risk exists (Pykhtin & Zhu, 2006)). While this can be seen as a limitation to our research, as e.g. Green (2015) and Rosen and Saunders (2012) report that right-/wrong way risk exists, we decide not to incorporate it as it would substantially increase the complexity and computation time of the exposure calculation. Under these assumptions, we can approximate the CVA at time t as follows

$$CVA(t) = \gamma \cdot \sum_{i=1}^{N+1} q_i(1 - q_i^o) \cdot EPE(t, T_i), \tag{5.1}$$

where we take γ to be the constant LGD for a counterparty and q_i the counterparty's market implied PD in the time period $(T_{i-1}, T_i]$ given that the counterparty did not default before this time period. q_i^o is the own market implied default probability in the period $(T_{i-1}, T_i]$ given that the default did not happen before this period, thus $1 - q_i^o$ is the corresponding own survival probability. The Expected Positive Exposure (EPE) discounted to time t is given by $EPE(t, T_i)$. It is calculated as follows

$$EPE(t, T_i) = \mathbb{E}^{\mathbb{Q}_{(1)}} \left[\tfrac{N(t)}{N(T_i)} \max(V(T_i), 0) \middle| \mathcal{F}_t \right], \quad t \leq T_i, \tag{5.2}$$

where we denote the filtration at time t as $\mathcal{F}_t$ and use that $\frac{N(t)}{N(T_i)}$ is the discount factor consisting of numéraires that is used to discount the future exposure to time t. Furthermore, we take the expectation with respect to the risk-neutral probability measure $\mathbb{Q}_{(1)}$ to exclude arbitrage opportunities. For a mathematical definition of the discrete-time numéraire as a function of zero-coupon euro bond prices, see Appendix A or Jamshidian (1997). Since $N(t)$ is known at time t, we can take it out of the expectation. Furthermore, the value of a portfolio at time T_i can also be defined as the expected discounted value of the future cash flows that follow from the portfolio, i.e.

$$V(T_i) = N(T_i) \mathbb{E}^{\mathbb{Q}_{(1)}} \left[\sum_{S \geq T_i} \frac{CF(S)}{N(S)} \middle| \mathcal{F}_{T_i} \right]. \tag{5.3}$$

Then, by rewriting $\max(V(T_i), 0)$ as $V(T_i)\mathbb{1}_{V(T_i)>0}$ and using the tower law of expectations and the expression for the portfolio value given in Equation (5.3), we can rewrite Equation (5.2) to

$$\begin{aligned}
EPE(t, T_i) &= N(t)\mathbb{E}^{\mathbb{Q}_{(1)}} \left(\frac{\max(V(T_i), 0)}{N(T_i)} \middle| \mathcal{F}_t \right), \\
&= N(t)\mathbb{E}^{\mathbb{Q}_{(1)}} \left(\frac{V(T_i) \cdot \mathbb{1}_{V(T_i)>0}}{N(T_i)} \middle| \mathcal{F}_t \right), \\
&= N(t)\mathbb{E}^{\mathbb{Q}_{(1)}} \left(N(T_i) \frac{\mathbb{E}^{\mathbb{Q}_{(1)}} \left[\sum_{S \geq T_i} \frac{CF(S)}{N(S)} \middle| \mathcal{F}(T_i) \right] \cdot \mathbb{1}_{V(T_i)>0}}{N(T_i)} \right), \\
&= N(t)\mathbb{E}^{\mathbb{Q}_{(1)}} \left(\mathbb{E}^{\mathbb{Q}_{(1)}} \left[\sum_{S \geq T_i} \frac{CF(S)}{N(S)} \middle| \mathcal{F}(T_i) \right] \cdot \mathbb{1}_{V(T_i)>0} \right), \\
&= N(t)\mathbb{E}^{\mathbb{Q}_{(1)}} \left(\mathbb{E}^{\mathbb{Q}_{(1)}} \left[\sum_{S \geq T_i} \frac{CF(S)}{N(S)} \mathbb{1}_{V(T_i)>0} \middle| \mathcal{F}(T_i) \right] \right), \\
&= N(t)\mathbb{E}^{\mathbb{Q}_{(1)}} \left(\sum_{S \geq T_i} \frac{CF(S)}{N(S)} \mathbb{1}_{V(T_i)>0} \right),
\end{aligned} \tag{5.4}$$

following Allick (2023). Here, we used the tower law of expectations between the penultimate and the last line. The future cash flows of the derivatives portfolio are straightforward to calculate, as the future cash flows of each derivative in the portfolio can be computed using contract details such as e.g. the strike price of the derivative and the simulated forward interest and forward exchange rates. Note that the discounted value of the future cash flows of a carbon forward can be calculated using Equation (3.1). Thus, we are left with estimating $\mathbb{1}_{V(T_i)>0}$, so the question whether the value of the derivatives portfolio is positive or not in case the counterparty defaults at time T_i. To this end, we use the Longstaff-Schwartz estimation method as explained in Longstaff and Schwartz (2001) and further elaborated in a CVA setting by Joshi and Kwon (2016). This is a Monte Carlo algorithm that estimates $V(T_i)$ based on a set of state variables. Assume that we generate K paths of forward and exchange rates from the CIRM model and take these rates together with a constant to be the set of b simulated state variables $\xi_{T_i,k}$ at time T_i in path k. Recall that the state variables are all interest and exchange rates in the CIRM model (see Section 3.1).

Let $U(T_i, k)$ denote the sum of future portfolio cash flows in the period $[T_i, T_{N+1}]$ that are discounted to time T_i in path k. Longstaff and Schwartz (2001) argue that this value can be approximated by the continuation value $f(T_i, k) = \sum_{j=1}^{b} \alpha_{T_i,j}\xi_{T_i,k,j}$, where $\alpha_{T_i,j}$ is the regression coefficient corresponding to $\xi_{T_i,k,j}$, which is state variable j at time T_i in path k. We use regression functions instead of the $U(T_i, k)$ directly as this allows us to potentially calculate exposure sensitivities with respect to the state variables, since the regression functions are differentiable while $U(T_i, k)$ is not. We estimate the regression coefficients for each tenor by performing a cross-path least squares regression (Joshi & Kwon, 2016) at each tenor date, thus for each tenor in $[t, T_{N+1}]$ we minimize

$$\sum_{k \in K} \left[\sum_{j=1}^{b} \alpha_{T_i,j}\xi_{T_i,k,j} - U(T_i, k) \right]^2 . \tag{5.5}$$

With the estimated coefficients and a set of simulated paths, we compute the continuation value $f(T_i, k)$ of the portfolio at time T_i. We then take this estimate as the Longstaff-Schwartz estimate of $V(T_i)$ as implied by the simulated state variables that follow from the CIRM model.

To calculate the PD of a counterparty, we decide to bootstrap these probabilities from market-observable credit default swap (CDS) quotes. This way, we obtain implied PDs for each period in the tenor structure based on market data. These probabilities may differ per counterparty and are treated as given. The LGD is the percentage of the notional value that is lost in case the counterparty defaults (Gregory, 2020). The LGD can also differ per counterparty. Similar to the PD, these can be extracted from market data and are thus treated as a known constant.

5.2 Debt Valuation Adjustment

Similarly, the DVA is the CVA from a counterparty's perspective, it can be seen as the fair value adjustment to account for own default risk. The concept of DVA has been somewhat controversial, as it is technically possible for a party to obtain (unrealized) profits due to their own credit deterioration (Madan & Schoutens, 2015). Nevertheless, DVA is important, as it allows a defaulting institution to claim (a part of) the remaining cash flows of their outstanding derivative contracts (Gregory & German, 2012). The discrete-time bilateral DVA is computed as follows

$$DVA(t) = \gamma^o \cdot \sum_{i=1}^{N+1} q_i^o (1 - q_i) \cdot ENE(t, T_i), \tag{5.6}$$

where γ^o is the constant own LGD, $1 - q_i$ the counterparty's survival probability given that no default happened before time T_i and $ENE(t, T_i)$ the Expected Negative Exposure (ENE). Similar as for CVA, the LGD, and PD are based on market data. Furthermore, the ENE is given by

$$ENE(t, T_i) = \mathbb{E}^{\mathbb{Q}_{(1)}} \left[\frac{N(t)}{N(T_i)} \min(V(T_i), 0) \middle| \mathcal{F}_t \right]. \tag{5.7}$$

Again, we use a Longstaff-Schwartz estimation method to compute $V(T_i)$ in Equation (5.7).

5.3 Netting and Collateral Agreements

One of the reasons why the positive exposure of a derivatives portfolio is unequal to the sum of the individual trade positive exposures is the presence of netting and/or collateral agreements. Netting agreements ensure that exposures of multiple contracts can be aggregated in case of a default (Pykhtin & Zhu, 2006). For example, suppose that an institution has one outstanding trade with a defaulting counterparty with a positive exposure of 2000 euros and one trade with a negative exposure of -1500 euros. In case of default without a netting agreement, the institution's positive exposure to the counterparty would be 2000 euros and needs to be paid to the counterparty. The money that the counterparty is still owed to the institution (the negative exposure on the second trade) will probably be only partially reimbursed. The remaining amount needs to be booked as a loss by the institution.

However, if a netting agreement is in place between the two parties, the netted exposure would be only $2000 - 1500 = 500$ euros at the time of default. Therefore, the positive exposure is significantly reduced by the fact that there is also a trade with a negative exposure in the portfolio. Although this is maybe a stylized example of how netting agreements work, the point is that netting agreements can reduce exposure for institutions. Note that netting agreements will never increase the positive exposure compared to a situation without netting agreements, hence the XVA will also always be lower than or equal to the XVA without netting. Due to its frequent usage by financial institutions, we allow for the existence of netting by incorporating it in the definition of $V(t)$ and by evaluating XVA per portfolio instead of per trade.

Similarly, a collateral agreement can also reduce the risk associated with a derivative contract. This is done through a collateral threshold. In short, if an institution imposes a collateral threshold on the counterparty, the counterparty needs to post collateral when the exposure of the contract exceeds the threshold. This way, the contract exposure is capped at the level of threshold, which reduces the potential exposure to the counterparty for the institution. It is also possible that the counterparty imposes a collateral threshold for the institution. Usually, collateral is paid using liquid assets like cash.

5.4 Analytical Benchmark Model

In order to assess the correctness of the CIRM model, we make use of the following benchmark model. This benchmark model is able to compute analytical approximations for the Expected Positive Exposure (EPE) and Expected Negative Exposure (ENE), and thus skips the complicated and stochastic simulation of interest and foreign exchange rates. From here onwards, we specifically discuss EPE calculations but ENE (and thus DVA) can be computed analogously. In short, this analytical benchmark model obtains EPE values by valuing multiple call options on carbon forward contracts using the Black formula. This way, an analytical EPE profile is obtained for the lifetime of the call. This EPE profile can then be used to compute the CVA as described before.

Recall from Section 5.1 that the EPE of a carbon forward can be computed as in Equation (5.2), hence as a discounted value of an expectation that includes the value of a carbon forward at the time of default. Additionally, in case we have a portfolio consisting of one carbon forward contract we can write the value of this portfolio as

$$V_{portfolio}(t) = B^{(1)}\left(t, T_i\right) \left(X^{(1:M)}\left(t, T_i\right) - K\right), \tag{5.8}$$

as follows from Equation (3.1) for the value of one single carbon forward at time t and expiring at time T_i. Before continuing, it is important to note that time T_i can correspond to any tenor in the tenor structure, and that we can split up the period between the value date (which is time $0 = T_0$) and time T_i in R periods of equal length. Additionally, note that we vary observation time t in order to obtain an entire exposure profile over the period $[0 = T_0, \ldots, T_i]$ rather than one EPE estimate. Specifically, we thus compute $EPE(t, T_i)$ at R different observation times t in $[0 = T_0, \ldots, T_i]$. Note that we do not need (all) t to correspond with a tenor date from the tenor structure, as we will see that we can compute $EPE(t, T_i)$ at any point in time.

To evaluate the EPE for time T_i analytically, we need to use a carbon forward contract that also expires at time T_i. Then, using the process $B^{(1)}\left(s, T_i\right)$ for s running in the interval $[0, t]$ as the numéraire, we can rewrite Equation (5.2) to

$$\begin{aligned} EPE(t, T_i) &= \mathbb{E}^{\mathbb{Q}_{(1)}}\left[\frac{B^{(1)}\left(0, T_i\right)}{B^{(1)}\left(t, T_i\right)}\max(V(t), 0)\right], \\ &= \mathbb{E}^{\mathbb{Q}_{(1)}}\left[\frac{B^{(1)}\left(0, T_i\right)}{B^{(1)}\left(t, T_i\right)}\max(B^{(1)}\left(t, T_i\right)\left(X^{(1:M)}\left(t, T_i\right) - K\right), 0)\right], \\ &= B^{(1)}\left(0, T_i\right)\mathbb{E}^{\mathbb{Q}_{(1)}}\left[\max(X^{(1:M)}\left(t, T_i\right) - K, 0)\right], \\ &= c\left(B^{(1)}\left(0, T_i\right), 0, t, X^{(1:M)}\left(0, T_i\right), K, \sigma_t^{(1:M)}\left(\cdot\right)\right), \end{aligned} \tag{5.9}$$

where $c(\cdot)$ denotes the Black formula of Black (1976) for a call option with as inputs the discount factor, current time, running time, forward rate, strike, and volatility. $\sigma_t^{(1:M)}\left(\cdot\right)$ is the volatility of the carbon forward exchange rate. In this analytical model we assume that this carbon volatility is constant and independent of maturity, hence no maturity is specified for the volatility parameter in Equation (5.9).

To go from the third to the fourth line, notice that the third line can also be interpreted as the counterparty risk-free value of a forward contract discounted to time 0 as noted by e.g. Arregui et al. (2017) and Salvador and Oosterlee (2021). This value can be computed using the well-known Black formula, as long as we assume that the carbon forward rate is lognormally distributed. Hence, we can analytically describe $EPE(t, T_i)$ and this can be done for any t in the interval $[0, T_i]$ as all parameters are known at any point in the period $[0, T_i]$. This allows us to obtain the EPE profile that we need to compute the CVA. All other steps in the CVA calculation are similar as for the simulated approach and are described in Section 5.1. Note that in Equation (5.9) we assume a bought position of an at-the-money carbon forward as we use the Black formula of a call. In case we want to use a sold position, we need to use the put version of the Black formula $p(\cdot)$. Note that for options that are at-the-money but otherwise identical, the Black formulas for calls and puts yield identical results.

There are several limitations to the benchmark model. First, in order to make use of the Black formula for each time t, we need to assume a certain volatility structure. Our assumption that the volatility of the carbon forward rate is independent of maturity is quite restrictive. Furthermore, the benchmark model cannot take the netting effects and/or collateral agreements of Section 5.3 into account whereas the simulated approach can. Lastly, although the analytical approach is quite straightforward to use when the portfolio consists of one derivative contract, it cannot value a portfolio of multiple contracts. In the case of a bigger portfolio, the benchmark model has to value each contract in the portfolio separately and thus fails to take any correlation between the payoffs of the contracts into account.

6 Results

In this section, we discuss the main findings of this paper. Similar to other papers on Valuation Adjustments (XVA) such as Joshi and Kwon (2016), Longstaff and Schwartz (2001) and Karlsson, Jain, and Oosterlee (2016), we assess the validity of the Commodity Interest Rate Market (CIRM) model and Valuation Adjustments (XVA) computations by studying them in real-life applications. One practical reason for this way of assessing performance is the lack of historical data, which prevents us from doing a relevant historical evaluation. In the analyses, we primarily make use of a one-year carbon forward contract that expires on 3 January 2025, has a carbon forward price of 77.69 euros for one carbon emission allowance, and consists of 5,000 allowances in total. Hence, the amount of money that needs to be transferred from the buyer to the seller at maturity equals 388,425 euros. In general, we analyze the situation in which we take a short position in the forward contract, as XVA are usually computed by large financial institutions and these institutions generally take a short position in a carbon forward contract. In principle, we evaluate the contract based on end-of-day data as of 3 January 2024 unless indicated otherwise. We get the end-of-day data through a commercial market data provider. Additionally, we assume that the (fictional) counterparty for the carbon forward has a known and constant Standard & Poor's credit rating. We assume that the LGD for the fictional counterparty is also known and fixed. Lastly, we only evaluate the validity of the methods in a setting where we only have two 'currencies' (the euro and carbon), but we would like to stress that the methods can also be applied in scenarios with more or other currencies.

In Section 6.1, we dive into the Mark-to-Market values that follow from the CIRM model for the carbon forward contract. Section 6.2 shows the behavior of the Expected Positive Exposure and Expected Negative Exposure when changing several characteristics of the carbon forward contract. Then, in Section 6.3 we disclose several results on the volatility parameters of the CIRM model and we present the results of several sanity checks on this. Lastly, in Section 6.4 we combine carbon forward contracts with other financial derivatives in one portfolio and test whether the models can still compute reasonable XVA. In this last section, we also investigate the effect of the correlation parameter between the euro interest rate and the carbon exchange rate.

6.1 Mark-to-Market Comparison

First, we compare the simulated carbon forward exchange rates of the CIRM model to the market-observed forward exchange rates. We do this by looking at the Mark-to-Market (MtM) values following from the CIRM and the benchmark model for the carbon forward contract. The MtM is defined as the fair value of a contract evaluated at current market prices. The MtM of a contract fluctuates over the lifetime of the contract, as market prices change frequently. In the case of a carbon forward contract, the value of the forward changes the most when there is a change in the carbon forward exchange rate. As Allen and Carletti (2008) indicate, using the MtM has benefits and limitations, as it reflects the true value of a contract, but it is also very volatile as it changes whenever market prices fluctuate. However, it can be a decent evaluation metric as it allows researchers to investigate whether the simulated carbon forward exchange rate curve that follows from the CIRM model matches the market-observed exchange rate curve.

We compute the MtM based on the market-observable forward exchange rate $X(t,T)$ for a bought carbon forward contract consisting of A carbon emission allowances, forward price K and expiry time T as

$$MtM_{analytical}(t,T) = A\left(X(t,T) - K\right), \tag{6.1}$$

where time T can correspond to any time (not necessarily corresponding to a tenor of the tenor structure as well). If the contractual forward price of the carbon forward is below the market-observed carbon forward price, we expect a MtM that is negative (below 0) as the buyer of the contract could have gotten allowances on the market for a lower price than it has to pay according to the contract at expiry. Alternatively, if the contractual forward price is above the market-observable forward exchange rate, we expect a positive MtM as the buyer has to pay less for its allowances than if he were to buy them on the market at expiry.

For the MtM based on the simulated carbon forward exchange rate of the CIRM model, we have to use an equivalent but slightly different formulation for the MtM. This is because the CIRM model simulates interest and exchange rates in each simulation path under the domestic spot measure and only for the tenor dates of the tenor structure. Then, following the first fundamental theorem of asset pricing, we compute the CIRM model-based MtM for the same bought carbon forward contract as

$$\frac{MtM_{CIRM}(t,T_i)}{N(t)} = \mathbb{E}^{\mathbb{Q}_{(1)}}\left[\frac{A(X(T_i)-K)}{N(T_i)}\right], \tag{6.2}$$

which holds only on the tenor dates of the tenor structure. Note that the definition of the numéraires in this equation are given in Appendix A, while $X(T_i)$ denotes the carbon spot exchange rate at time T_i. In theory, both expressions for the MtM have the same meaning. The formulas slightly differ due to the setup of the CIRM model, where we have a simulated forward exchange rate in each path and thus need to take the expectation over all paths to obtain the MtM. As the $X(t,T)$ term in Equation (6.1) is only one exchange rate observed at time t, there is no need to take an expectation.

The forward price in the carbon forward contract (77.69 euros per allowance) is chosen such that the market-observed carbon forward exchange rate curve is matched exactly at contract inception. This means that the initial analytical MtM at inception is equal to 0. We would expect that the simulated MtM is equal to zero as well. However, it might be possible that differences are observed due to simulation or discretization errors since Equation (6.2) takes an expectation with respect to the domestic spot measure over all simulation paths.

When evaluating the MtMs of the short carbon forward contract at three different evaluation moments, we get the results as shown in Table 1. The table also gives the absolute difference between the MtMs of both models. We find that the MtM at inception is indeed 0.00 for the analytical model, while the MtM of the simulated model is 278.24 euros at contract inception. This can be explained by the fact that we used the market-observed carbon forward exchange rate curve as provided by the data provider for the determination of the forward price, while the simulated forward exchange rate curve can be slightly different from this curve. In practice, this simulation error is relatively small compared to the notional of the contract. Since there are 5,000 allowances in the contract, we find that the difference in carbon forward exchange rate equals $278.24/5,000 = 0.06$ cents at inception.

To show that the differences between the simulated and market-observed carbon forward exchange rate curves are not of considerable size, we compute a one-sample t-test with the null hypothesis that the means are equal. We can use this test since we have independent simulation paths. Furthermore, this test can be performed as the independent paths together are (approximately) normally distributed and because we have many simulation paths (5,000 to be exact). As the results show, we cannot reject the null hypothesis at any significance level, all p-values are quite high. Hence, we conclude that the simulated and analytical MtMs do not significantly differ from one another.

Table 1. Mark-to-Market Values for a 1 Year Carbon Forward Contract

Evaluation Date	Analytical MtM	Simulated MtM	Absolute Difference	p-values
3 January 2024	0.00	278.24	278.24	0.91
10 January 2024	37,107.37	37,310.51	203.14	0.93
17 January 2024	69,934.51	70,042.14	107.63	0.96

Note. This table shows the Mark-to-Market (MtM) results in euros for a short carbon forward contract that expires on 3 January 2025. The table gives the MtMs following from the simulated approach based on the CIRM model, the MtMs following from the analytical approach that uses the Black formula, and the absolute difference between the two MtMs. The MtMs are computed at three evaluation moments. At each evaluation date, we also compute a p-value to indicate whether the difference in MtM is statistically significant.

When evaluating the MtMs one week and two weeks later, we see that both models find that the carbon forward exchange rate curve has changed. The MtMs are on both dates a bit larger for the simulated model than for the analytical one, which is no surprise as it also simulated a slightly lower forward exchange rate at inception. We observe that the absolute differences between the MtMs are small, ranging from 203.14 on 10 January to 107.63 on January 17th. Moreover, we see that the p-values in Table 1 become even higher than for 3 January. Based on these findings and the finding in Appendix C on the MtM for a contract with a different maturity date, we believe that the CIRM model can simulate forward exchange and forward interest rates that are in line with the rates observed in the market.

To investigate the convergence of the simulated model to the analytical model when increasing the number of simulation paths, we show the MtMs for a varying number of paths in Figure 3. For our results in Table 1, we used 5,000 simulation paths. As mentioned before, in each path we simulate spot forward exchange rates, euro forward interest rates, and convenience forward rates. Therefore, in each path we have a simulated carbon forward exchange rate, so we compute a MtM in each path. Then, in case we have 5,000 paths, we take the average over all paths and obtain the MtM as shown in Table 1.

Figure 3. Convergence of the CIRM Model to the Benchmark Model

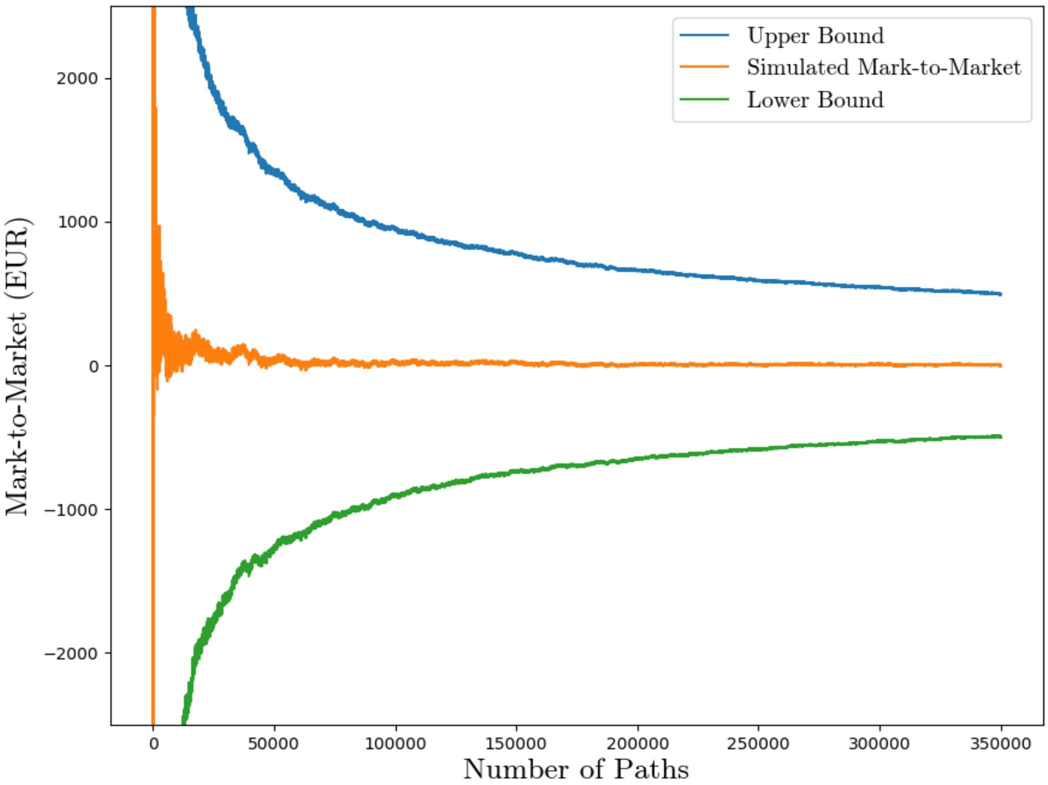

Note. Figure shows the evolution of the CIRM model-based Mark-to-Market (MtM) in euros for our carbon forward contract when we increase the number of paths that are used to compute the MtM. The MtMs are computed for 3 January 2024. Furthermore, we see that the simulated MtM converges to the analytical MtM when we increase the number of paths as the analytical MtM has a value of zero on 3 January 2024. We also show the upper- and lower bounds of the 90% confidence intervals.

We find that when we increase the number of used paths for the simulation, we get a MtM that is closer to the analytical MtM of 0 at inception. The simulated MtM essentially converges to the analytical MtM. Figure 3 shows that using more than 50,000 paths does not lead to large movements in the simulated MtM anymore. To show that the simulated MtM does not significantly differ from the analytical one, we plot the upper- and lower bounds of the 90% confidence intervals of the simulated MtM. Figure 3 also shows that the analytical MtM indeed falls in the 90% confidence interval at every point in the figure, which is another indication that the simulated MtM does not significantly differ from the analytical one. Based on the results of this section, we decide to use 5,000 simulation paths for the other analyses, as this gives MtMs that are not significantly different from the analytical model and this number ensures that we have the lowest computation time possible. Namely, when using 50,000 or more paths, the computation time of the CIRM model increases substantially, with little significant improvement compared to using 5,000 paths.

6.2 Exposure Results

Next, we dive into the Expected Positive Exposure (EPE) and Expected Negative Exposure (ENE). First, we give the exposure results in Figure 4. As expected, we see that the EPE as computed by both the CIRM and the benchmark model is always larger than or equal to zero, while the ENE cannot be positive. Furthermore, we see that the differences between the analytical and simulated EPE and ENE are small, indicating that the exposures based on the analytical model are closely matched by the ones following from the CIRM model. Figure 4a also shows that the observed differences in initial MtM have a negligible effect on the total EPE and ENE, as the exposures from both models are also similar in the first month of the exposure profiles.

Another observation from Figure 4a is that the EPE (ENE) exposure is increasing (decreasing) as time progresses. This is no coincidence and we explain the financial intuition behind this behavior for the EPE profile. Recall that the initial MtM of the carbon forward contract is zero and that the exposure depends on the interest and exchange rates of the CIRM model. These rates are all simulated/constructed at the valuation time, which in Figure 4a corresponds to the time of contract inception. From financial theory, we know that volatility increases with time because there is an increasing probability that the state variables will be further away from the initial levels as time increases. As the value of the contract depends on the level of the state variables, we have that the contract value will also be further away from its initial (zero) value as time progresses. However, as Equation (5.2) highlights, in the EPE calculations we only consider the paths where the contract value at the time of default is positive. Therefore, the paths with a negative contract value are set to zero. As the positive paths get on average more extreme as we get further away from the valuation time, we see that the total exposure also increases as time progresses.

We also check what the exposure profiles look like when computing the exposures almost one month later on 1 February 2024 in Figure 4b. As observed before, we see that the EPE (ENE) is increasing (decreasing) as time progresses. Moreover, we find that the EPE does not start at zero anymore, which makes sense as the MtM of the contract is also unequal to zero on 1 February. The MtM of the contract has changed as the carbon spot exchange rate has decreased substantially since 3 January. Due to the changes in the carbon spot price, we see that the EPE profile starts at roughly 75,000 euros, which leads to higher exposures than originally estimated at contract inception. It is important to regularly monitor the exposure on a contract, as it indicates the open risk position that the seller of the contract has to the buyer. If the seller believes its open risk position to the buyer is too large, they can decide to e.g. potentially unwind or refinance the trade before its time to maturity is reached.

On the other hand, the ENE is higher (less negative) than on 3 January. Since the ENE can only attain values below zero, the ENE starts at zero and only slowly decreases over time, leading to a higher ENE overall than at contract inception. Additionally, the ENE decreases less over time than in Figure 4a. This is because the initial MtM is positive instead of zero, due to which fewer exposure paths attain negative exposure levels and even if these negative exposure levels are reached, they are less extreme than in the zero initial MtM case. Due to the initial MtM at 75,000 euros, the exposure paths are essentially skewed upwards which leads to higher EPE and lower ENE values.

Figure 4. Simulated and Analytical Expected Exposure Paths

(a) Evaluated at 3 January 2024 **(b) Evaluated at 1 February 2024**

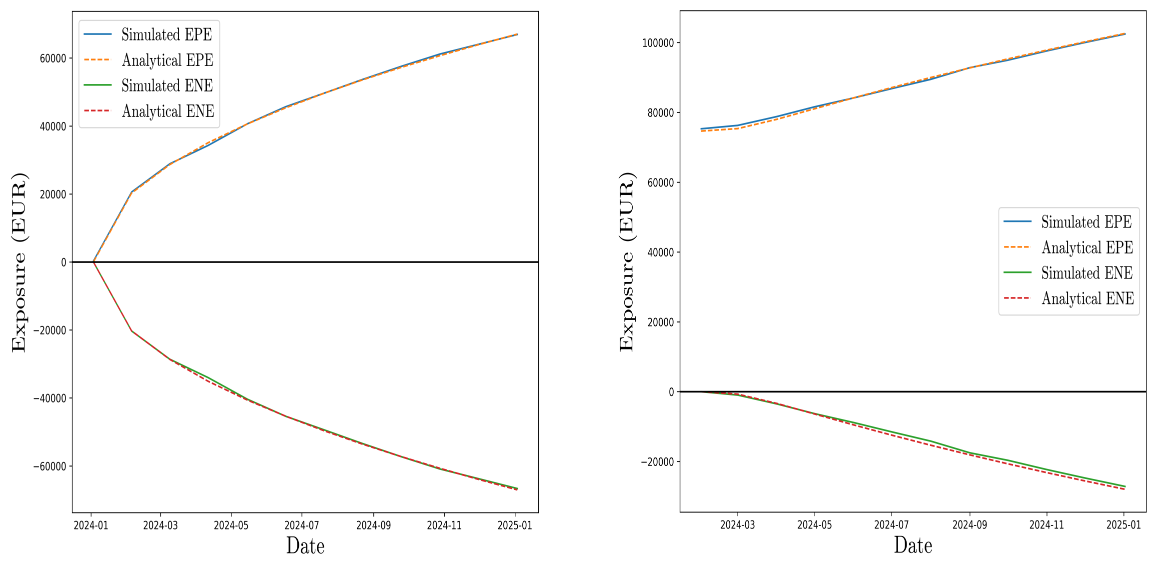

Note. Figure shows the Expected Positive Exposure and Expected Negative Exposure in euros over the lifetime of the one-year carbon forward expiring on 3 January 2025. In this contract, we take a short position. The exposures are plotted both for the analytical benchmark model as well as the CIRM model. The left figure shows the exposure profiles evaluated on 3 January 2024, so at contract inception. The right figure shows the exposure profiles on 1 February 2024.

As one of the advantages of the CIRM model over the benchmark model is that we can also compute exposures for collateralized carbon forward contracts, we study the exposures of a collateralized contract in Figure 5. In this figure, we evaluate the exposures for the same one-year carbon forward contract, but now we impose a collateral threshold of 60,000 euros both for the institution and the counterparty. As a reference, we also plot the EPE and ENE profiles when there are no collateral agreements in place, these profiles correspond to the profiles for the simulated exposures in Figure 4a. We find that imposing a collateral threshold substantially reduces the EPE and ENE. We explain this finding for the EPE but the ENE reduction can be explained analogously. Due to the threshold of 60,000 euros, the positive exposures are effectively capped at this level in each simulation path. Therefore, paths that originally attained exposure levels that were higher than 60,000 euros are set to this level which reduces the total EPE of the contract. However, not every path originally attained a positive exposure level of 60,000. These paths are left unchanged due to the collateral agreement, and their lower levels ensure that the total EPE is generally not near the threshold level but maximally reaches a level of 20,000-30,000 euros. We therefore find that the exposures following from the CIRM model can successfully take collateral agreements into account, as they indeed reduce the EPE and ENE to a level substantially below the threshold.

Figure 5. Expected Exposure Paths with Collateral Agreements

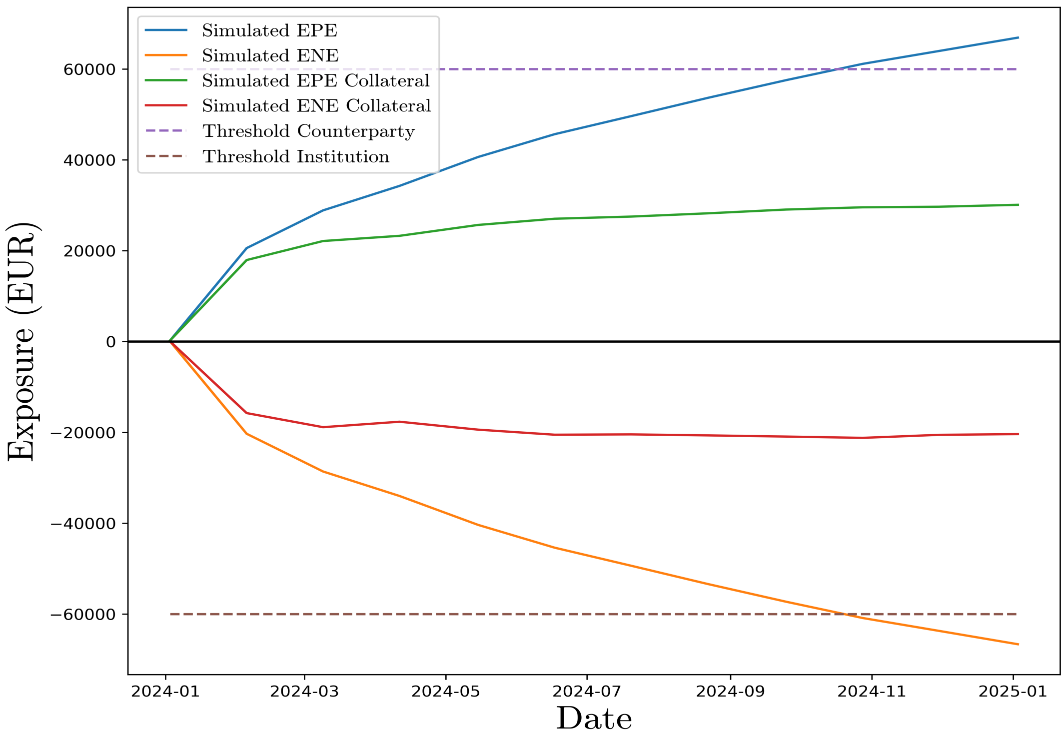

Note. Figure shows the Expected Positive Exposure (EPE) and Expected Negative Exposure (ENE) in euros when there are two collateral agreements in place for a one-year carbon forward contract expiring on 3 January 2025. The exposures are plotted over the lifetime of the forward contract. Two collateral agreements are in place, one for the institution and one for the counterparty. Both thresholds are set at a level of 60,000 euros.

Lastly, we perform several robustness checks for the EPE profiles by evaluating carbon forward contracts with different notionals and different times to maturity. As the results for these contracts are comparable to what we described before, we show the results of these robustness checks in Appendix C.

6.3 Impact of the Carbon Volatility Parameter

In this section, we show the effects of changing the volatility parameters in the CIRM model. To show the effects of the interest and exchange rate volatility parameters in practice, we plot 2,000 positive exposure (PE) paths that are used to calculate the EPE for the short one-year carbon forward contract in Figure 6. Elsewhere in this research, the EPE is determined based on 5,000 paths, but due to readability reasons we choose to only show the first 2,000 paths. We also show the total simulated EPE that follows from the 5,000 paths by a red dashed line, which is the same line as the line for the simulated EPE in Figure 4a. In Figure 6, we see that there is quite some dispersion between the 2,000 paths in terms of PE, as values attain levels higher than 200,000 euros and lower than -600,000 euros. Obviously, the bigger the dispersion between the different paths, the more variability there is and the larger the volatility parameters are.

The finding that quite some paths attain 'positive' exposures below zero euros might come as a surprise. An explanation for this lies in Equation (5.4). Note that the EPE is, by definition, nonnegative as shown by Equation (5.2). However, as we see in the last line of Equation (5.4), the PE estimate can be below zero when the cash flow of the forward contract after exposure evaluation time is negative, while the contract overall has a positive value when the exposure is evaluated. To make this more specific, a negative cash flow at maturity can happen when (in the case of a sold forward contract) the forward price of the contract is below the at contract maturity prevailing carbon spot price. However, at the exposure evaluation time, the contract can have a positive MtM value regardless of the carbon spot price at contract maturity and therefore the positive exposure path is not set to zero for the EPE calculation. Therefore, we also see some paths with negative exposure levels, but by looking at the density of the paths in Figure 6, there are far more positive exposure paths that have a positive MtM at exposure evaluation time and a positive cash flow at contract maturity.

Figure 6. Positive Exposure Paths for a Carbon Forward Contract

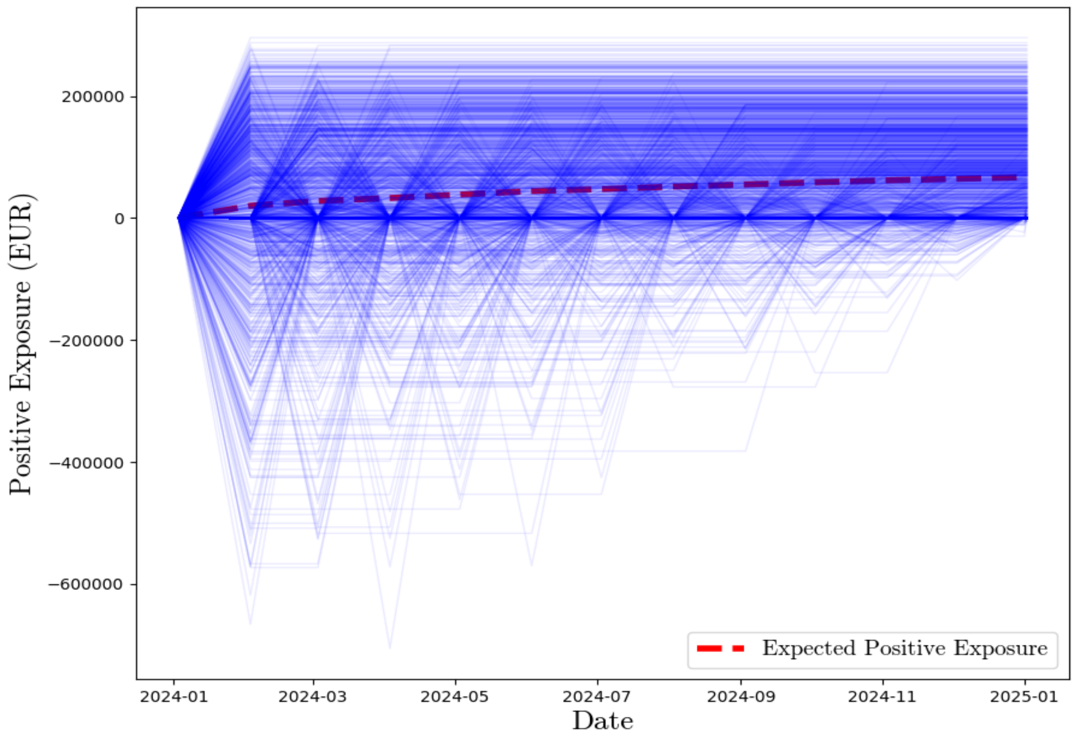

Note. Figure shows 2,000 exposure paths that are used to calculate the total Expected Positive Exposure in euros of the short carbon forward contract. Note that the entire simulation model simulates 5,000 paths, but for readability reasons we choose to only show 2,000 paths. The red dashed line shows the Expected Positive Exposure that follows from all simulated paths. Recall that if we take the average of all the blue lines at each point in time, we would obtain the red line. The positive exposures are denoted in euros.

We dive into the effects of changing the carbon exchange rate volatility parameter in Figure 7. In blue, we see the EPE as also shown in Figure 4a. When the carbon volatility parameter is significantly increased, we see that the EPE of the carbon forward contract increases as well. An increase in carbon volatility could be due to e.g. increased uncertainty in the (carbon emission allowance) market. The increase in EPE due to a ceteris paribus increase in carbon volatility can be explained as follows. With a larger volatility parameter, we get more extreme simulated carbon forward exchange rates in each path compared to a situation with lower volatility. Therefore, the forward contract valuations in all simulation paths will also be more excessive than in the normal volatility case. However, when we compute the EPE, we get that all paths with negative forward contract values will be set to 0, as a party does not have a positive exposure to a contract if it has a negative value for them at the time of their counterparty's default. Nonetheless, the more extreme paths with positive portfolio values will be taken into account when computing the EPE. Therefore, we end up with an overall EPE that is higher than in a setting with lower volatility.

Figure 7. Expected Positive Exposure with Different Volatility Parameters

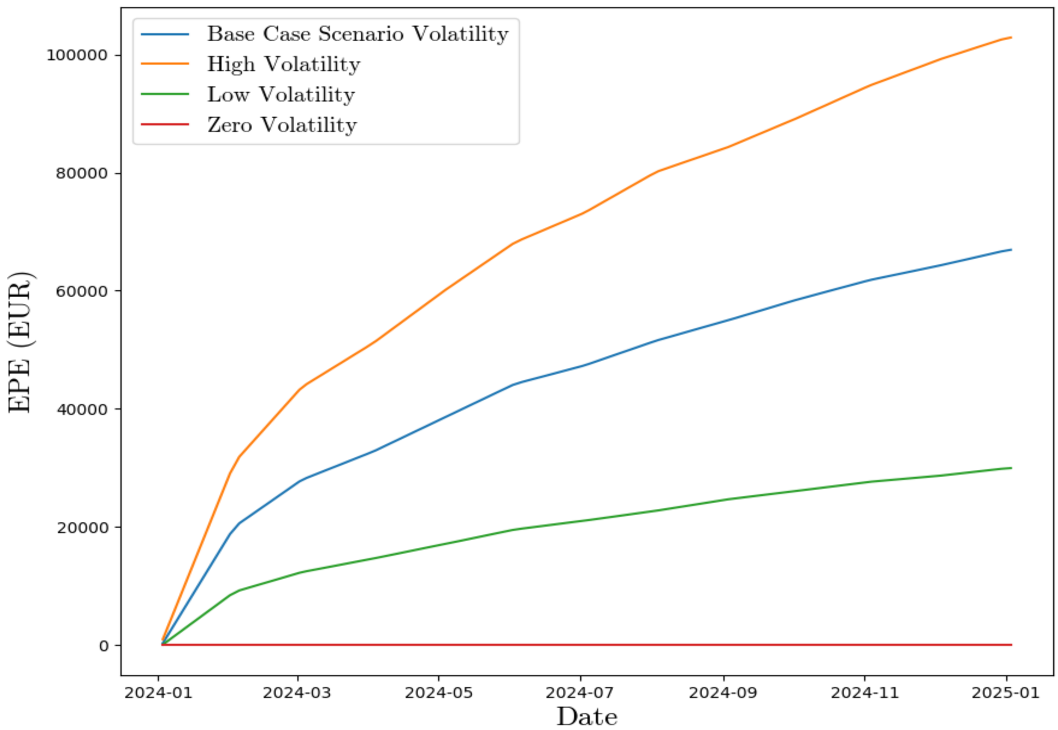

Note. Figure shows the Expected Positive Exposure (EPE) in euros when we change the carbon exchange rate volatility. We plot EPEs using different levels of carbon forward exchange rate volatility for the Marked-to-Market carbon forward contract and we evaluate the EPEs using data as of 3 January 2024 (contract inception).

Similarly, if we set the carbon exchange rate volatility lower than in the base case scenario, we obtain less extreme simulated carbon forward exchange rates which, in turn, lead to less excessive carbon forward values. Again, since paths in which the value is negative are set to zero, we end up with an EPE estimate based on the less dispersed paths and have a lower EPE than in the base case volatility scenario. We also investigate the scenario in which we have zero carbon price volatility so that the carbon exchange rate does not change at all over the lifetime of the contract. This results in a constant EPE at the level of the MtM at contract inception. As the carbon forward is at-the-money at contract inception, we have that the EPE is always equal to zero.

6.4 Portfolio and XVA Results

In this last subsection, we dive into the exposure profiles of portfolios that consist of multiple derivative contracts. Specifically, we consider the case where we have a portfolio of the aforementioned carbon forward contract combined with a receiving position in a euro interest rate swap. We evaluate this scenario as this can occur frequently for European financial institutions that offer carbon forward contracts to other companies. In short, an interest rate swap is a contract between two parties in which they agree to exchange a series of interest payments without paying off the underlying debt (Bicksler & Chen, 1986). Imagine a situation where a company has a one-year euro loan with a bank for which it has to pay a variable interest rate and monthly coupons. In case the company wants to get rid of its interest rate risk that arises due to the variable (also called floating) interest rate, it can enter an interest rate swap with another counterparty in which it pays a fixed interest rate to the counterparty, while the counterparty will pay a floating interest rate to the company. This way, the company only has to pay a fixed rate and can use the proceedings from the swap to pay its floating rate obligations on the euro loan.

Usually, a financial institution would take a receiving position in a euro interest rate swap. This means that they pay a floating interest rate to their corporate client, while they receive a fixed rate from their counterparty. The usual EPE profile for such a euro swap would roughly take the shape of a concave parabola. This is because yield/interest rate curves are usually upward-sloping so that the floating rate in the swap is below the fixed rate at first. However, as time progresses the floating rate starts to increase to the level of the fixed rate since the fixed rate remains constant. This behavior causes the swap to first experience an increase in exposure, as the contract first has a positive value for the institution since they make a net profit in the beginning. This profit arises from the fact that the floating rate (which the institution pays) is below the fixed rate (which the counterparty receives). However, when the floating rate starts to match the fixed rate, the top of the exposure profile is reached, and when the floating rate starts to increase beyond the level of the fixed rate, the exposure decreases again as now the institution must pay more to the counterparty than they receive from them. This decrease in exposure continues until contract maturity, as then the exposure is equal to zero.

Figure 8 shows the EPE profiles of several derivative portfolios which include carbon forwards. We combine the previously mentioned carbon forward with a one-year 30 million euro interest rate swap with monthly coupons, and in this contract a receiving position is considered (so that the floating rate is paid and the fixed rate is received). If we combine both derivatives in one portfolio and calculate the portfolio's EPE, we obtain the blue line in Figure 8. This blue exposure path takes potential correlation and netting effects for the portfolio into account. When we evaluate the EPEs of both derivatives separately and sum these EPEs, we get the red exposure path in the figure. It is no surprise that this way of computing EPEs yields the highest exposures. This is because the exposure of an entire portfolio will always be smaller than or equal to the exposure of the two products separately, or in mathematical form

$$(Exposure_{forward} + Exposure_{swap})^+ \leq (Exposure_{forward})^+ + (Exposure_{swap})^+, \qquad (6.3)$$

where the $(\cdot)^+$ function is zero or higher depending on whether the term in brackets is larger than zero. To see the rationale behind Equation (6.3), think of the following example. Assume that the exposure to the forward contract is equal to -20 euros and +10 euros to the swap in one specific path. Then, the left-hand side of Equation (6.3) will be equal to $-20 + 10 = -10$ euros, hence this side will be equal to zero. The right-hand side, however, will be equal to $0 + 10 = 10$ euros, hence Equation (6.3) holds. Similarly, in case both contracts have positive or negative exposure in a path, the left- and right-hand sides of Equation (6.3) will be equal. Therefore, our finding that adding exposures separately yields the highest EPE intuitively makes sense.

Figure 8. Expected Positive Exposure for Portfolios with Carbon Forwards

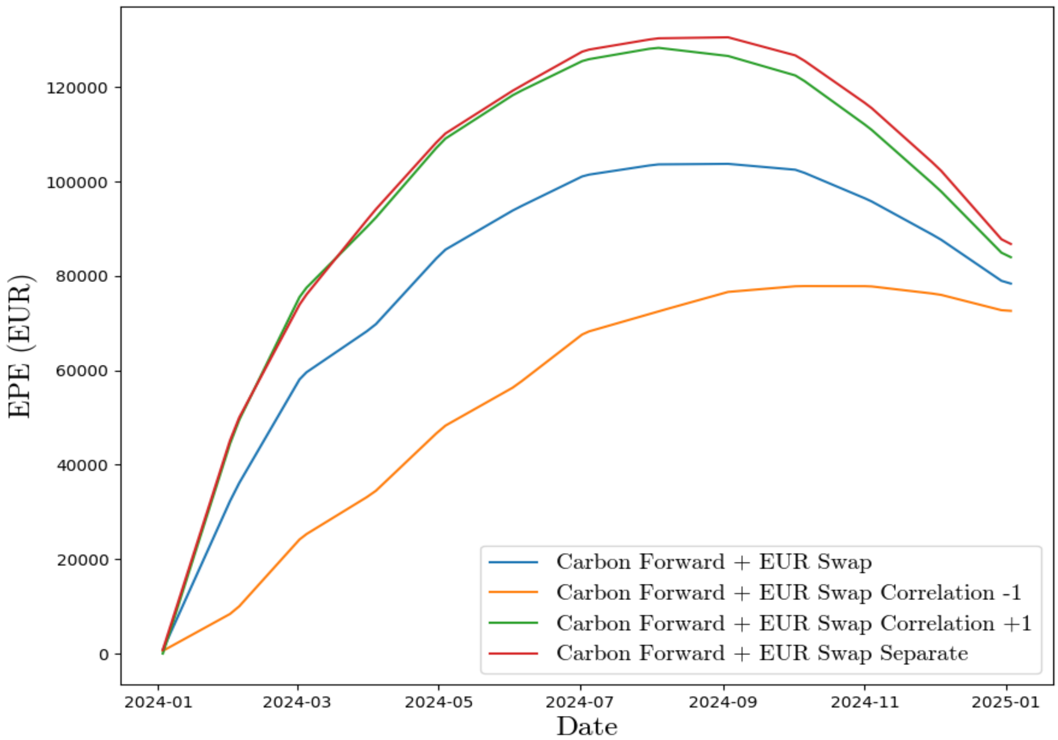

Note. Figure shows the Expected Positive Exposure (EPE) for a portfolio consisting of a carbon forward and a one-year euro interest rate swap with different underlying correlations. Additionally, we plot the EPEs when we take a carbon forward and a euro swap together in one portfolio and we show the sum of the two contract exposures when evaluated separately. The EPE is denoted in euros.

Additionally, we evaluate the effects of altering the correlation between the euro interest rate and the carbon exchange rate in Figure 8. We find that a correlation parameter of +1 (perfect positive correlation) gives higher exposures for the portfolio of a sold carbon forward and a receiving euro swap, while a correlation of -1 (perfect negative correlation) gives the lowest EPE. As expected, the EPE profile with an estimated correlation parameter is in between the ones for perfect positive and negative correlation.

Our finding on the correlation effects can be explained by using economic rationale. In case of perfect negative correlation, an increase in the euro interest rate causes an equal percentage decrease in the carbon price. As a consequence, we obtain a lower exposure for the receiving swap in the portfolio, since we pay the (increased) floating rate in the swap contract and thus the swap becomes less valuable.

At the same time, the decrease in carbon price causes the sold carbon forward to become more valuable. This is because allowances need to be delivered at maturity and this can be done by buying them at a lower price on the market, thus making the contract itself worth more money. Hence, the exposure related to this contract goes up. We therefore see that the changes in exposure of both contracts are in opposite direction, which causes the positive exposures to be less extreme, and thus causing a lower overall EPE.

In case there is perfect positive correlation, we observe the opposite effect. If the euro interest rate goes down, then the carbon price decreases as well. The downfall in the euro interest rate causes the swap to become more valuable, thus increasing exposure on the swap, while the decrease in the carbon price increases the value of the forward contract, thus increasing the exposure on this. Hence, we see that exposures on both contracts increase in this scenario, thus making the positive exposure paths more extreme, which increases the portfolio's EPE.

Lastly, we dive into the XVA results for several derivative portfolios in Table 2. These results are in line with our findings on the portfolio exposures, as we again see that the XVA for a portfolio of a forward with a swap are significantly lower than when summing the individual XVA. Note that the DVA is always less than or equal to zero, as the Probability to Default and the Loss Given Default are always positive but the ENE is always less than or equal to zero. The DVA can thus reduce the total XVA, which makes sense as the DVA is related to the own default risk and can therefore be seen as a benefit to the counterparty. Based on the findings that the analytical benchmark model renders EPE and ENE profiles that are very similar to the simulation-based exposures, we conclude that we can accurately calculate XVA for carbon forwards using the CIRM model.

Table 2. XVA Results

Portfolio	CVA	DVA
Carbon forward	117.22	-67.81
Euro swap	138.60	-18.46
Carbon forward + EUR swap	207.15	-57.07

Note. This table shows the Valuation Adjustments (XVA) for three different derivative portfolios. All XVA are computed as of 3 January 2024 and are denoted in euros.

7 Conclusion

The purpose of this paper is to calculate Valuation Adjustments (XVA) for carbon forwards in the framework of a Commodity Interest Rate Market (CIRM) model. To achieve this goal, we calibrate a CIRM model using market-observed prices on several financial products and simulate forward interest and forward exchange rates from this calibrated market model. Based on these simulated rates, we compute the Expected Positive Exposure (EPE) and Expected Negative Exposure (ENE) for a portfolio of derivative contracts and individual carbon forwards. Then, we combine the EPE and ENE results with the Probability to Default (PD) and the Loss Given Default (LGD) to obtain XVA for derivative portfolios.

We evaluate the validity of the CIRM model using an analytical benchmark model and economic rationale. We start by checking whether the simulated carbon forward exchange rate curve of the CIRM model matches the market-observed forward exchange rate curve. This is done by looking at the Mark-to-Market (MtM) values based on the CIRM and the benchmark model. We find that the MtMs from the CIRM model are close to the ones for the benchmark model (which uses the market-observed carbon exchange rate curve for its MtM calculation), and the differences between the MtM values are not statistically significant. We therefore conclude that the simulated carbon forward exchange rates are an accurate representation of the market-observed carbon forward exchange rates. We also investigate whether we can improve the CIRM model to match the market-observed carbon exchange rate curve even more, and it turns out that we can. This can be achieved by increasing the number of simulation paths. If we increase the number of paths to 200,000 simulation paths, we get MtMs that are basically identical. However, this comes at a cost of increased computation time, while the accuracy gains are minor.

Additionally, we determine whether the exposure profiles of the simulated model coincide with the profiles of the benchmark model. We observe that no clear distinction can be made between the two models, thus giving another indication that the simulated rates from the CIRM model are closely matching the interest and exchange rates in the market. Furthermore, as a sanity check we also evaluate the EPE profiles when we manually increase the volatility of the carbon forward exchange rate, and we find that this leads to an upward shift in EPE.

Moreover, we investigate whether we are indeed able to construct reasonable EPE profiles for collateralized carbon forwards and a portfolio of derivative contracts. These cases are of particular importance, as the CIRM model's ability to compute exposures in these cases is one of its biggest benefits compared to the benchmark model and other approaches. It turns out that we can compute reasonable EPE profiles for derivative portfolios and collateralized trades as well. This indicates that the CIRM model can successfully take netting and collateralization effects into account.

Consistent with our economic intuition, we find that a portfolio of a short carbon forward with a receiving euro interest rate swap experiences increased EPE when we have perfect correlation between the euro interest rate and the carbon forward exchange rate. Similarly, negative correlation leads to the lowest possible EPE for this portfolio. Finally, we compute XVA and see that incorporating the Debt Valuation Adjustment (DVA) decreases the total XVA. All in all, these empirical findings lead us to the conclusion that XVA can be accurately calculated using the simulated interest and exchange rates of the CIRM model, making the CIRM model a successful extension of the market model of Pilz and Schlögl (2013) with a practical application in XVA modeling.

Despite the conclusion that XVA can be accurately computed using the CIRM model, we do have some suggestions for further improvements of the proposed methods. First of all, we suggest allowing for a volatility smile for the interest rate volatilities in the CIRM model. In this research, we only use at-the-money caps and floors to calibrate these volatilities, due to which we have a flat implied volatility curve. In practice, implied volatility curves exhibit some skew and curvature (as shown in e.g. Deuskar et al. (2008)) which causes options that are far in- or out-of-the money to have a higher implied volatility. However, with Black implied interest rate volatilities, we are not able to model this smile due to the assumptions of the model. As a consequence, it is possible that the exposures that follow from the CIRM model are underestimated since we model lower volatilities than with a volatility smile. A potentially underestimated exposure could lead to underestimated XVA, which is undesirable from a risk standpoint. A possible alternative to Black implied volatilities that does allow for a volatility smile would be using (piecewise constant) volatility skew functions as proposed by Karlsson, Pilz, and Schlögl (2016) which builds upon the calibration proposed by Piterbarg (2005). The results of Karlsson, Pilz, and Schlögl (2016) look promising, but their methods require the usage of option data on carbon futures, which is only sparsely available and trading in these options is not very liquid. If data on these options becomes more freely available, calibrating a volatility smile could be something worthwhile to look into.

Another limitation of this research is as well related to the general lack of liquid (option) data for carbon. Specifically, in this research we perform a historical analysis for the calibration of the correlation parameters in the CIRM model. This implies a calibration method under the real-world measure, while the stochastic differential equations of the CIRM model are all under the risk-neutral measure. A remedy to this miscalibration would be to use data on so-called spread options on interest and exchange rates. Unfortunately, data on these kinds of options has limited availability, or in the case of carbon, does not exist. As the results show, changing the correlation parameter has a substantial influence on the EPE profile, hence if data on spread options becomes more accessible it is recommended to alter the currently implemented correlation calibration procedure.

Lastly, we suggest incorporating wrong-way risk in the XVA calculation. When allowing for wrong-way risk we allow the exposure to the counterparty to be dependent on the counterparty's probability to default in the XVA calculation. It is well documented that wrong-way risk exists (see e.g. Green (2015) and Rosen and Saunders (2012)) and that accounting for it would generally lead to a substantial increase in XVA (Hull & White, 2012). Although relatively easy rules of thumb exist for incorporating wrong-way risk, more sophisticated approaches would include the usage of copulas (see e.g. Brigo and Chourdakis (2009) and Cespedes et al. (2010)) or the modeling of the default probability as a function of exposure (Hull & White, 2012). However, the downside of these methods is their increased complexity and longer computation time for the XVA calculation. Nevertheless, incorporating wrong-way risk would make the XVA calculation even more accurate, thus making future research into this topic a valuable addition to the literature on XVA modeling.

References

Allen, F., & Carletti, E. (2008). Mark-to-market accounting and liquidity pricing. *Journal of Accounting and Economics*, *45*(2-3), 358–378.

Allick, D. (2023). *LIBOR Market Model calibration for credit valuation adjustment of a portfolio of swaptions.* Retrieved from http://hdl.handle.net/2105/65835.

Areal, N., Rodrigues, A., & Armada, M. R. (2008). On improving the least squares Monte Carlo option valuation method. *Review of Derivatives Research*, *11*, 119–151.

Arregui, I., Salvador, B., & Vázquez, C. (2017). PDE models and numerical methods for total value adjustment in European and American options with counterparty risk. *Applied Mathematics and Computation*, *308*, 31–53.

Baldeaux, J., & Platen, E. (2015). Credit derivative evaluation and CVA under the benchmark approach. *Asia-Pacific Financial Markets*, *22*, 305–331.

Benz, E., & Trück, S. (2009). Modeling the price dynamics of CO2 emission allowances. *Energy Economics*, *31*(1), 4–15.

Bicksler, J., & Chen, A. H. (1986). An economic analysis of interest rate swaps. *The Journal of Finance*, *41*(3), 645–655.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, *3*(1-2), 167–179.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654.

Boyle, P. P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics*, *4*(3), 323–338.

Brace, A., Gatarek, D., & Musiela, M. (1997). The market model of interest rate dynamics. *Mathematical Finance*, *7*(2), 127–155.

Brigo, D., & Chourdakis, K. (2009). Counterparty risk for credit default swaps: Impact of spread volatility and default correlation. *International Journal of Theoretical and Applied Finance*, *12*(07), 1007–1026.

Brigo, D., Chourdakis, K., & Bakkar, I. (2009). Counterparty risk valuation for energy-commodities swaps: Impact of volatilities and correlation. *arXiv preprint arXiv:0901.1099*.

Brigo, D., & Mercurio, F. (2006). *Interest rate models-theory and practice: With smile, inflation and credit* (Vol. 2). Springer.

Broadie, M., Glasserman, P., & Jain, G. (1997). Enhanced Monte Carlo estimates for American option prices. *Journal of Derivatives*, *5*, 25–44.

Carriere, J. F. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics*, *19*(1), 19–30.

Cespedes, J. C. G., de Juan Herrero, J. A., Rosen, D., & Saunders, D. (2010). Effective modeling of wrong way risk, counterparty credit risk capital, and alpha in Basel II. *The journal of risk Model Validation*, *4*(1), 71.

Chang, S., & Wang, X. (2015). Modelling and computation in the valuation of carbon derivatives with stochastic convenience yields. *Plos One*, *10*(5), 1–21.

Cox, J. C., Ingersoll Jr, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, *53*, 769–799.

Daskalakis, G., Psychoyios, D., & Markellos, R. N. (2009). Modeling CO2 emission allowance prices and derivatives: Evidence from the European trading scheme. *Journal of Banking & Finance*, *33*(7), 1230–1241.

Deuskar, P., Gupta, A., & Subrahmanyam, M. (2008). The economic determinants of interest rate option smiles. *Journal of Banking & Finance*, *32*(5), 714–728.

European Commission. (2023). *What is the EU ETS?* Retrieved from https://climate.ec.europa.eu/eu-action/eu-emissions-trading-system-eu-ets/what-eu-ets_en.

Fernández, J. L., Pou, M., & Vázquez, C. (2015). A drift-free simulation method for pricing commodity derivatives. *Applied Stochastic Models in Business and Industry*, *31*(4), 536–550.

Gamba, A. (2003). Real options valuation: A Monte Carlo approach. *Faculty of Management, University of Calgary WP*, (2002/3).

Garman, M. B., & Kohlhagen, S. W. (1983). Foreign currency option values. *Journal of International Money and Finance*, *2*(3), 231–237.

Green, A. (2015). *XVA: Credit, funding and capital valuation adjustments*. John Wiley & Sons.

Gregory, J. (2020). *The XVA challenge: Counterparty risk, funding, collateral, capital and initial margin*. John Wiley & Sons.

Gregory, J., & German, I. (2012). Closing out the DVA debate. *Available at SSRN 2790204*.

Härle, P., Lüders, E., Pepanides, T., Pfetsch, S., Poppensieker, T., & Stegemann, U. (2010). Basel III and European banking: Its impact, how banks might respond, and the challenges of implementation. *EMEA Banking*, 16–17.

Heath, D., Jarrow, R., & Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society*, 77–105.

Ho, T. S., & Lee, S.-B. (1986). Term structure movements and pricing interest rate contingent claims. *the Journal of Finance*, *41*(5), 1011–1029.

Hull, J., & White, A. (1990). Pricing interest-rate-derivative securities. *The Review of Financial Studies*, *3*(4), 573–592.

Hull, J., & White, A. (2012). CVA and wrong-way risk. *Financial Analysts Journal*, *68*(5), 58–69.

Jamshidian, F. (1997). LIBOR and swap market models and measures. *Finance and Stochastics*, *1*, 293–330.

Joshi, M., & Kwon, O. K. (2016). Least squares Monte Carlo credit value adjustment with small and unidirectional bias. *International Journal of Theoretical and Applied Finance*, *19*(08), 1650048.

Joy, C., Boyle, P. P., & Tan, K. S. (1996). Quasi-Monte Carlo methods in numerical finance. *Management Science*, *42*(6), 926–938.

Karlsson, P., Jain, S., & Oosterlee, C. W. (2016). Counterparty credit exposures for interest rate derivatives using the stochastic grid bundling method. *Applied Mathematical Finance*, *23*(3), 175–196.

Karlsson, P., Pilz, K. F., & Schlögl, E. (2016). Calibrating a market model to commodity and interest rate risk. *Available at SSRN 2773974*.

Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies*, *14*(1), 113–147.

Lovcha, Y., Perez-Laborda, A., & Sikora, I. (2022). The determinants of CO2 prices in the EU emission trading system. *Applied Energy*, *305*, 117903.

Madan, D. B., & Schoutens, W. (2015). Conic CVA and DVA. *Robert H. Smith School Research Paper No. RHS*, *2715403*.

Miltersen, K. R., Sandmann, K., & Sondermann, D. (1997). Closed form solutions for term structure derivatives with log-normal interest rates. *The Journal of Finance*, *52*(1), 409–430.

O'Connor, F. A., Lucey, B. M., Batten, J. A., & Baur, D. G. (2015). The financial economics of gold—A survey. *International Review of Financial Analysis*, *41*, 186–205.

Pearson, K. (1895). VII. Note on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, *58*(347-352), 240–242.

Pilz, K. F., & Schlögl, E. (2013). A hybrid commodity and interest rate market model. *Quantitative Finance*, *13*(4), 543–560.

Piterbarg, V. (2005). Time to smile. *Risk*, 71–75.

Pykhtin, M., & Rosen, D. (2010). Pricing counterparty risk at the trade level and CVA allocations. *FEDS Working Paper*.

Pykhtin, M., & Zhu, S. H. (2006). Measuring counterparty credit risk for trading products under Basel II. *RISK Books*.

Rogers, L. C. (2002). Monte Carlo valuation of American options. *Mathematical Finance*, *12*(3), 271–286.

Rosen, D., & Saunders, D. (2012). CVA the wrong way. *Journal of Risk Management in Financial Institutions*, *5*(3), 252–272.

Rostagno, M., Altavilla, C., Carboni, G., Lemke, W., Motto, R., & Guilhem, A. S. (2021). Combining negative rates, forward guidance and asset purchases: Identification and impacts of the ECB's unconventional policies.

Salvador, B., & Oosterlee, C. W. (2021). Total value adjustment for a stochastic volatility model. a comparison with the Black–Scholes model. *Applied Mathematics and Computation*, *391*, 125489.

Schlögl, E. (2002). A multicurrency extension of the lognormal interest rate market models. *Finance and Stochastics*, *6*, 173–196.

Seifert, J., Uhrig-Homburg, M., & Wagner, M. (2008). Dynamic behavior of CO2 spot prices. *Journal of Environmental Economics and Management*, *56*(2), 180–194.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, *5*(2), 177–188.

World Bank. (2023). *State and trends of carbon pricing 2023.* Retrieved from http://hdl.handle.net/10986/39796n.

Zhu, S. H., & Lomibao, D. (2005). A conditional valuation approach for path-dependent instruments. *Available at SSRN 806704*.

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

Appendix

A Numéraires

In this section, we explain the concept of numéraires. A numéraire is defined as a unit of currency against which we can measure other currencies. One famous example of a numéraire is the US dollar in the market for oil, as all oil prices are expressed in US dollars. We denote the numéraire for time t by $N(t)$. In this research, we take the discretely compounded euro money market account $N(t)$ as the numéraire at time t. With this numéraire, we can discount future cashflows from, for example, time T to time t by using the discount factor $\frac{N(t)}{N(T)}$. Note that the numéraire is uniquely defined on the tenor dates only (Jamshidian, 1997), and that we do not interpolate the numéraire values between tenor dates.

The idea behind the money market account as numéraire is that one could invest 1 euro at time T_0 and buy a zero-coupon bond that expires at time T_1. Then, one could reinvest the proceeds at time T_1 in a zero-coupon bond that expires at time T_2, and so on. This way, the numéraire process based on the discrete euro money market account is formally given by

$$\begin{aligned}
N(0) &= 1, \\
N(1) &= \frac{1}{B^{(1)}\left(T_0, T_1\right)} B^{(1)}\left(T_1, T_2\right), \\
N(2) &= \frac{1}{B^{(1)}\left(T_0, T_1\right)} \frac{1}{B^{(1)}\left(T_1, T_2\right)}, \\
&\vdots \\
N(t) &= B^{(1)}\left(t, T_{i+1}\right) \prod_{j=0}^{i-1} \frac{1}{B^{(1)}\left(T_j, T_{j+1}\right)},
\end{aligned} \tag{A.1}$$

hence the numéraire is determined by the prices of the zero-coupon bonds denominated in euros. With the numéraire process, we can e.g. calculate the risk-neutral present value of a derivatives portfolio.

B Euler discretization

Here, we briefly explain how the stochastic differential equations (SDEs) can be used to simulate the spot forward exchange rates of the CIRM model. To translate SDEs into actual rates, we apply a logarithmic Euler scheme to discretize the rates. Consider an arbitrary SDE that is given by

$$dS_t = \mu(S_t, t)dt + \sigma(S_t, t)dW_t, \tag{B.1}$$

where S_t is the rate that needs to be simulated and W_t denotes a Wiener process observed at time t. If we integrate dS_t over the (small) period dt, we get

$$S_{t+dt} = S_t + \int_t^{t+dt} \mu(S_u, u)du + \int_t^{t+dt} \sigma(S_u, u)dW_u. \tag{B.2}$$

We can approximate both integrals with the left-point rule. When doing so, we obtain

$$\begin{aligned}\int_t^{t+dt} \mu(S_u, u)du &\approx \mu(S_t, t) \int_t^{t+dt} du, \\ &= \mu(S_t, t)dt,\end{aligned} \tag{B.3}$$

and

$$\begin{aligned}\int_t^{t+dt} \sigma(S_u, u)dW_u &\approx \sigma(S_t, t) \int_t^{t+dt} dW_u, \\ &= \sigma(S_t, t)(W_{t+dt} - W_t), \\ &= \sigma(S_t, t)\sqrt{dt}Z,\end{aligned} \tag{B.4}$$

where Z is a standard normal variable. Using these approximations in Equation (B.2) gives us the desired Euler discretization scheme

$$S_{t+dt} = S_t + \mu(S_t, t)dt + \sigma(S_t, t)\sqrt{dt}Z. \tag{B.5}$$

However, since we have lognormally distributed state variables, we need to use the logarithmic Euler discretization scheme. To obtain this, we apply the invertible transformation $f(Y_t) = e^{Y_t}$. Then, by using Itô's lemma we get

$$dY_t = \left(\frac{\mu(S_t, t)}{S_t} - \frac{\sigma^2(S_t, t)}{2S_t^2}\right) dt + \frac{\sigma(S_t, t)}{S_t} dW_t, \tag{B.6}$$

where we used that $S_t = e^{Y_t}$. Lastly, if we write the standard Euler scheme for Y_t and apply the aforementioned transformation, we obtain the logarithmic Euler scheme for S_t as given by

$$S_{t+dt} = S_t \cdot \exp\left(\left(\frac{\mu(S_t, t)}{S_t} - \frac{\sigma^2(S_t, t)}{2S_t^2}\right) dt + \frac{\sigma(S_t, t)}{S_t} dtZ\right), \tag{B.7}$$

where Z again denotes a variable that is drawn from a standard normal distribution.

C Robustness Checks with Other Carbon Forwards

This appendix shows several robustness checks that we perform to assess whether the results are consistently achieved in line with our expectations. First, we show the Expected Positive Exposures (EPE) and Expected Negative Exposures (ENE) for a carbon forward contract that has a higher notional than the carbon forward in Section 6. Specifically, we have increased the notional by doubling the number of allowances in the contract, so now we have 10,000 allowances instead of 5,000. The strike price of the contract remains the same so the contract has a Mark-to-Market (MtM) of zero at contract inception on 3 January 2025. The results are shown in Figure 9. In this figure, we observe similar patterns for all exposures as found previously in Section 6.2. Also, the EPE and ENE based on the CIRM model again closely match the exposure profiles of the analytical benchmark model, thus indicating that our conclusions also hold for contracts with a higher notional. However, we also see that the exposures are significantly larger than found before. This is no surprise, as a higher notional of the contract generally leads to a larger potential exposure to a counterparty.

Figure 9. Expected Exposure Paths for a Larger Carbon Forward Contract

(a) Evaluated at 3 January 2024

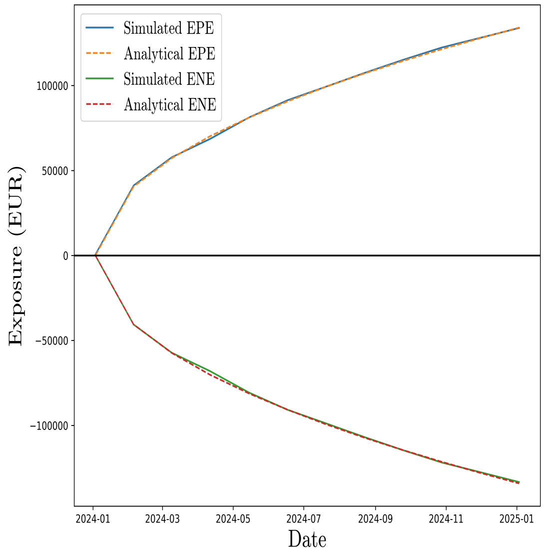

(b) Evaluated at 1 February 2024

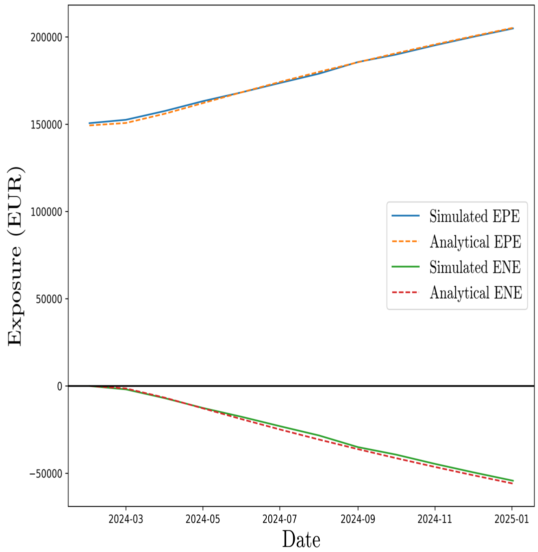

Note. Figure shows the Expected Positive Exposure and Expected Negative Exposure in euros over the lifetime of a one-year carbon forward expiring on 3 January 2025 with a notional that is twice as large than for the contract evaluated in Section 6. In this contract we take a short position. The exposures are plotted both for the analytical benchmark model as well as the CIRM model. The left figure shows the exposure profiles evaluated on 3 January 2024, so at contract inception. The right figure shows the exposure profiles on 1 February 2024.

Second, we check whether the CIRM model can also compute exposures in case we take a long position in the same carbon forward contract as evaluated in Section 6. The resulting exposure profiles are shown in Figure 10. We find that the EPEs and the ENEs of the CIRM model are again very close to the exposures following from the analytical model. Since the EPEs and ENEs evaluated at 3 January in Figure 4a are equally large in absolute terms, we find that the exposures in Figure 10a are also symmetric and identical to the ones obtained before. Differences occur when evaluating the exposures 1 month later in Figure 10b. Since we now take an opposite position in the contract, we get that the EPE of Figure 10b is equal to the ENE of Figure 4b multiplied by minus one. Similarly, the ENE in Figure 10b is equal to the EPE of Figure 4b multiplied by minus one.

Figure 10. Expected Exposure Paths for a Bought Carbon Forward

(a) Evaluated at 3 January 2024

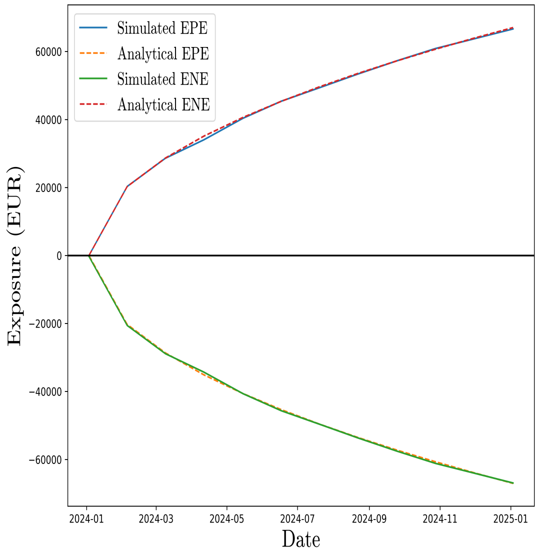

(b) Evaluated at 1 February 2024

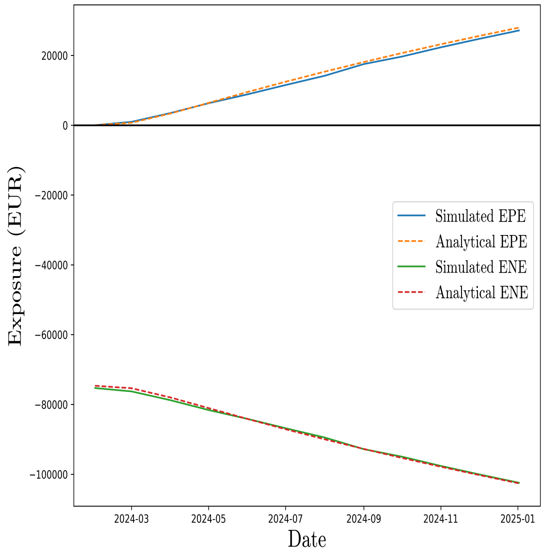

Note. Figure shows the Expected Positive Exposure and Expected Negative Exposure in euros over the lifetime of a one-year carbon forward expiring on 3 January 2025 with the same notional as for the contract evaluated in Section 6. In this contract we take a buy position. The exposures are plotted both for the analytical benchmark model as well as the CIRM model. The left figure shows the exposure profiles evaluated on 3 January 2024, so at contract inception. The right figure shows the exposure profiles on 1 February 2024.

As a third robustness check, we assess the performance of the CIRM model compared to the analytical benchmark model using a short carbon forward contract that has a time to maturity of five years. The exposure profiles for this contract are shown in Figure 11. The contract again consists of 5,000 allowances but has a different strike price than the one-year short carbon forward contract evaluated in the main text, as the time to maturity is different. However, we still ensure that the contract is marked-to-market at contract inception. We see that the simulated model has a slightly harder time matching the benchmark exposures, but no systematic under- or overestimation of the exposures is observed. As expected, we still see that the EPE (ENE) is steadily increasing (decreasing), and that more extreme exposure levels are reached compared to the one-year forward contract, which is no surprise as the longer time to maturity requires us to simulate forward interest and forward exchange rates for a longer period in the future, due to which we observe more extreme values. As mentioned before, more extreme (simulated) rates lead to more extreme exposure profiles.

Figure 11. Expected Exposure Paths for a 5-Year Carbon Forward Contract

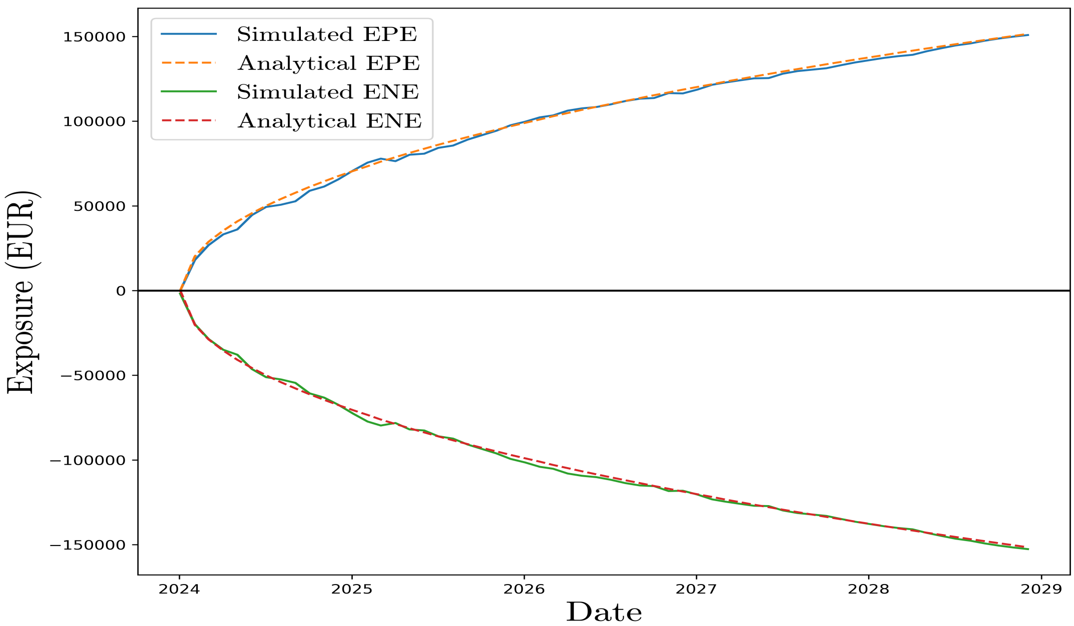

Note. Figure shows the Expected Positive Exposure and Expected Negative Exposure in euros on 3 January 2024 over the lifetime of a five-year carbon forward expiring at 3 January 2029. The contract is marked-to-market at contract inception and consists of 5,000 allowances. In this contract we take a sell position. The exposures are plotted both for the benchmark model as well as the CIRM model.

Besides exposure paths, we also perform a robustness check on the Mark-to-Market (MtM) values. Specifically, we evaluate the MtM of the three aforementioned carbon forward contracts with different specifications in Table 3. As expected, we obtain identical results for the bought carbon forward contract as for the sold contract in the main text, but now the simulated MtM is negative instead of positive as we take a different position in the contract. Furthermore, we see that the contract with a twice as large notional obtains a twice as large initial difference in MtM. As also the standard error is twice as large, we find identical p-values of 0.91. Lastly, we see that the short contract with a longer maturity has a larger error in initial MtM than the one-year contract, indicating that as time to maturity increases, the CIRM model has a harder time matching the market-observed carbon forward exchange rate curve than for shorter-dated maturities. However, note that the error relative to the notional of the contract is still small, and carbon forwards with a time to maturity of five years are not often observed in practice.

Table 3. Mark-to-Market Values for Three Different Carbon Forward Contracts

Forward Contract	Analytical MtM	Simulated MtM	Absolute Difference	p-values
High Notional	0.00	556.48	556.48	0.91
Bought	0.00	-278.24	278.24	0.91
Longer Maturity	0.00	-1761.37	1761.37	0.82

Note. This table shows the Mark-to-Market (MtM) results in euros for three different carbon forward contracts. The table gives the MtMs following from the simulated approach based on the CIRM model, the MtMs following from the analytical approach that uses the Black formula, and the absolute difference between the two MtMs. The MtMs are computed at contract inception for all three contracts. For each contract, we also compute a p-value to indicate whether the difference in MtM is statistically significant.