Potential Future Exposure in the Presence of Initial Margin

James Nevin

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in partial fulfilment of the requirements for the degree of Master of Philosophy.

August 1, 2019

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Philosophy to the University of Cape Town. It has not before been submitted for any degree or examination.

Signed by candidate

James Nevin

August 1, 2019

Abstract

This dissertation considers the concept of potential future exposure, and how initial margin can be used to mitigate it. In addition to this, the cost of implementing initial margin is estimated, and some of the difficulties associated with it are addressed.

The two primary techniques for calculating initial margin considered are nested Monte Carlo, and Gaussian Least Squares Monte Carlo. These two techniques are compared for effectiveness. It is shown that the nested Monte Carlo technique performs well under numerous conditions, and that the Gaussian Least Squares Monte Carlo relies on particular model and instrument characteristics.

Acknowledgements

I would like to thank my supervisor, Thomas McWalter, for his continued time, patience and encouragement. I owe an enormous amount to his ideas and support.

I would also like to thank David Taylor, Glen Point Capital, and the rest of the AIFMRM staff for making this journey possible.

I would also like to thank my parents for their financial and emotional support through all my studies, and my sister for making me laugh.

Lastly, I would like to thank God, for opening the doors of opportunity and leading me through every day of my life.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Prior to the financial crisis of 2008, there was far more limited regulation on financial institutions than is in place today. Often, modelling and pricing was implemented with an implicit assumption of little to no risk of counterparties defaulting. This was most apparent in the misuse of collateralised debt obligations in the United States, where the pooling of risk was treated as sufficiently safe. Post-crisis, regulation has made it crucial to acknowledge and objectively value default risk. One such measure of this risk is potential future exposure (PFE).

Counterparty exposure is defined as the value that would be lost should a counterparty default at a particular point in time — this value is the greater of zero and the market value of the portfolio (Canabarro and Duffie, 2003). Potential future exposure considers the maximum counterparty exposure at some point in the future, for a given confidence level. Often, PFE is mapped over a time horizon — for example, it may track the exposures over the period of one particular derivative, from inception to maturity. Expected exposure is similar to PFE, in that it measures average counterparty exposure at a given point in time.

Since PFE is included in the Basel III regulatory framework for measuring counterparty credit risk (Basel Committee and others, 2014), any company wishing to comply with Basel III needs to model and implement safety measures for PFE. This is often done through collateralisation. The goal of collateralisation is to reduce risk and to keep exposure relatively stable from inception to maturity. However, this is not always achieved in practice, as can be seen in Figure 1.1. In this example, the PFE over time modelled by Monte Carlo simulation for an FX call option is shown. It has the following characteristics and assumptions: the 99^{th} quantile of portfolios is used for PFE; exchange rates follow standard geometric Brownian motion (GBM), and deterministic interest rates; notional of USD (foreign currency) one million with maturity of one year; spot price for USD of 13 ZAR (domestic currency); interest rates r_{ZAR} of 8%, r_{USD} of 1.5%; volatility of stochastic process, σ_t, of 30%; in-the-money strike of 11.5 ZAR, out-of-the-money strike of 16 ZAR. The

collateral posted at any time is the value of the option at that time.

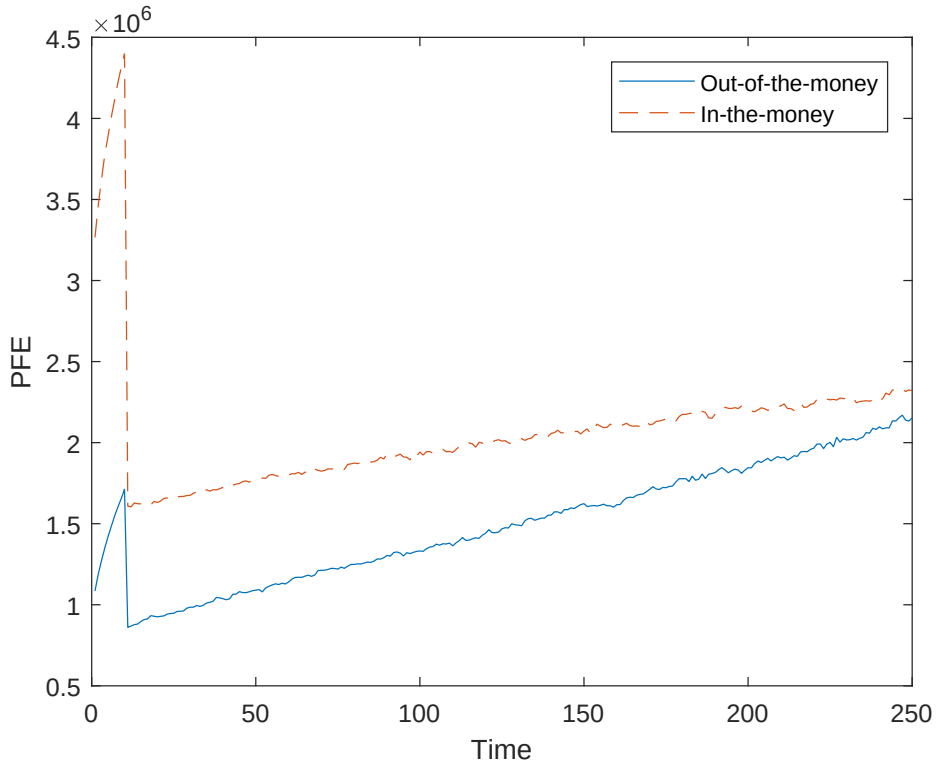

Fig. 1.1: PFE of FX call option for different strike prices.

The initial spike in both cases is as a result as what is known as the margin period of risk (MPOR), which is the time taken to close-out a position once a counterparty has defaulted. We assume that there is a time lag between entering into the contract and the posting of the first margin. Thus, the spike comes into effect when the counterparty does not post any margin at all. With effective margining, one might hope that the PFE decreases to zero following the MPOR. However, this is not the case in this example. This illustrates the necessity of more rigorous collateralisation techniques.

This dissertation aims to cover the calculation, effectiveness, and cost of initial margin (IM), a form of collateralisation which is intended to reduce the risk in over-the-counter trades. It will be shown that IM is sufficient for eliminating PFE, but there are some caveats.

In the next chapter, all of the relevant terminology will be introduced. In Chapter 3, various techniques for calculating initial margin will be detailed. Chapter 4 shows the results of the techniques in the context of an FX call option, as demonstrated above. Chapter 5 tests the primary techniques on a fixed-for-floating inter-

est rate swap. Finally, Chapter 6 summarises the results.

Chapter 2

Literature Review

2.1 Counterparty Credit Risk

In exchange-traded contracts, there is very low credit risk, since the exchange itself becomes the counterparty for the derivative. As a result, the focus of counterparty risk is on over-the-counter (OTC) transactions. However, margins still need to be computed for central counterparties.

Pykhtin and Zhu (2007) provide a basic outline for the modelling of credit risk exposure and pricing of counterparty risk. In the case where there is no counterparty risk mitigation, they explain that the maximum loss for a bank or financial institution when multiple contracts are held with a single, defaulting counterparty is given by:

$$E(t) = \sum_i E_i(t) = \sum_i \max\{V_i(t), 0\}, \tag{2.1}$$

where $E_i(t)$ and $V_i(t)$ are, respectively, the exposure and value on the i^{th} contract at time t. They suggest that this exposure can be limited through the use of netting agreements, which aggregates contracts between the parties in the event of default. This results in the following exposure:

$$E(t) = \max\left(\sum_i V_i(t), 0\right). \tag{2.2}$$

In order for financial institutions to manage their portfolios, it is necessary for them to model their exposure. Zhu and Pykhtin identify three main components for modelling the distribution of credit exposure: scenario generation, where future market scenarios are simulated at a discrete set of time points, $t_0, t_1, \ldots, t_n$; instrument valuation, where the value of each $V_i(t_j)$ is calculated at each time point; and portfolio aggregation, where $E(t_j)$ is calculated as per (2.1) or (2.2), depending on whether a netting agreement is in place.

Scenario generation can either be performed by creating a path from each t_j to the next, or by independently generating scenarios at each t_j. Which one is used depends on the properties of the derivatives in the portfolio — for example, sample paths will be necessary for path-dependent derivatives like American options.

The scenarios are usually specified by stochastic differential equations (SDEs). The example shown in Chapter 1 defines the scenario using GBM. Should the solution to the SDE being considered be known, the scenarios can be simulated directly. If not, approximations such as the Milstein method are necessary.

Once scenarios have been generated, the value of the instruments can be calculated. Zhu and Pykhtin recommend against computationally expensive techniques such as Monte Carlo, since it is often necessary to evaluate at a number of time points. Lomibao and Zhu (2005) showed that many of the conditional expectations of future instrument values given a set of scenarios have closed-form solutions. In opposition to this, Canabarro and Duffie (2003) identify Monte Carlo as a core part of the systems currently used to compute PFE. In the case of valuing less standard derivatives, Monte Carlo may prove necessary.

Performing aggregation using the various instrument valuations, PFE can be constructed using the portfolio values generated through the scenario simulations. The simulations will create a distribution for portfolio values at each of the discrete time points considered, and this distribution can be used to estimate high percentile confidence levels for future exposure. Zhu and Pykhtin identify two main factors influencing future exposure: diffusion and amortisation. As time progresses, there is a higher likelihood the scenarios diverge from the current state of the world, leading to higher future exposure — a diffusion effect. This is offset by amortisation, if the portfolio consists of numerous derivatives or cashflows of varying maturity, since there are fewer derivatives on which the counterparty can default over time.

2.1.1 Margined Portfolios

With the increased regulation and awareness of the risk associated with counterparty default, it has become common for financial institutions to include margin agreements in their OTC portfolios. These generally require one or both parties to post collateral when the exposure, $E(t)$, exceeds a certain threshold. Usually, if the exposure drops below the threshold, the collateral can be returned in part or in full. These arrangements aim to limit the growth of possible losses due to default.

Associated with the posting of collateral is a margin period of risk (MPOR). Andersen, Pykhtin and Sokol (2016) define the MPOR as the time period after the posting of collateral during which the gap between the portfolio value and collateral posted can widen. This is as a result of "a sequence of contractual and opera-

tional time lags", which take place after the counterparty has defaulted and before the position can be closed-out.

In the example in Chapter 1, the collateral at t is set to the mark-to-market value of the collateral arrangement at time t. This illustrates a category of collateralisation, known as variation margin (VM), in which the difference between portfolio value and collateral posted is kept below a certain threshold (Andersen *et al.*, 2016). However, due to delays in either closing-out a position once a counterparty has defaulted or identifying that a counterparty has defaulted, this collateral can only be accounted for once the MPOR has passed. As such, the collateralised exposure at time t is given by:

$$\hat{E}(t) = \max\{E(t) - E(t - \text{MPOR}), 0\},$$

where $\hat{E}(t)$ is the exposure of the collateralised portfolio at time t.

In order to address the discrepancies that can arise during the MPOR, an initial margin (IM) can be posted at each margining date, in addition to variation margin. Initial margin takes account of possible changes in the value of the portfolio over the margin period of risk. The posted IM is usually equal to a certain confidence level for changes in portfolio value; for example, a counterparty may be required to post IM equal to the 99% confidence level for portfolio changes over the MPOR. It is important to note that IM will often be required by both parties in a trade. The holder of an in-the-money contingent claim need not place any variation margin, but would be required to post initial margin if the 99% confidence level for changes in portfolio value moves the claim value below zero. In the case where both parties are required to post IM, netting is not permitted between them (Andersen, Pykhtin and Sokol, 2017a).

Finally, it is important to note that IM and VM are posted on each margining date, and that the time difference between margining dates does not necessarily need to correspond to the MPOR.

2.2 Valuation Adjustments

Valuation adjustment (XVA) is the practice of adjusting the fair value of contingent claims by taking into account additional costs and risks. This process is necessary for the allotment of capital required by Basel III, and for hedging. There are numerous types of XVA: credit valuation adjustment (CVA) is one of the oldest forms of XVA, where risk of counterparty default is accounted for; debt valuation adjustment (DVA) measures the risk of the financial institution itself defaulting; funding valuation adjustment (FVA) is used when there is a mismatch between the value of

the claim and the collateral offered. In this case, there is a need for additional hedging or capital allocation during the lifetime of the derivative. The costs associated with implementing these in the future are known as FVA; margin valuation adjustment (MVA) tracks the cost of implementing an initial margin for a derivative; capital valuation adjustment (KVA) allows for the costs associated with allocating capital to derivatives for regulatory purposes.

Andersen, Duffie and Song (2019) identify controversy surrounding the calculation and implementation of FVA. Implementation of FVA can conflict with other valuation adjustments, such as DVA. Additionally, a number of inconsistencies are present between practitioners in how FVA is calculated. These issues make objectively valuing FVA challenging.

In this dissertation, the calculation of MVA will be addressed for the various techniques of calculating initial margin considered. As part of the analysis comparing various techniques, the MVA values will be compared. Difficulties in calculating MVA will also be discussed.

Chapter 3

Methodology

Three values will be calculated and considered in various cases: variation and initial margins; potential future exposure, and how it is affected by VM and IM; and the cost of implementing initial margin, termed margin value adjustment. These values will be compared between the techniques explained below.

3.1 Variation and Initial Margins

In all cases, VM is calculated on a simulation-by-simulation basis. We only consider contingent claims and financial models that result in closed-form solutions for the value of the claim. This is necessary, as we will be required to value claims at various periods for each simulation.

3.1.1 Calculation of Initial Margin

Suppose that there exists a filtered probability space with a multivariate set of state variables X_t, which are adapted to the filtration at time t. These state variables are the variables that influence the value of the contract which we wish to work with. For example, for the case given in Chapter 1, the only state variable is the exchange rate. Two techniques are presented here for calculating initial margin under these conditions: nested Monte Carlo, and a technique termed Gaussian Least Squares Monte Carlo (GLSMC) — this terminology was introduced by McWalter, Kienitz, Nowaczyk, Rudd and Acar (2018).

Nested Monte Carlo

Using a nested Monte Carlo algorithm with n outer simulations and m inner simulations, one is able to obtain unbiased estimates for initial margin for a given contingent claim at time t using the following procedure:

1. Jointly simulate n values for the multivariate state process X_t

2. For each of the n simulated X_t values, simulate m values for $X_{t+\text{MPOR}}$
3. For each of the n X_t values, set the initial margin to the 99% quantile of the m $X_{t+\text{MPOR}}$ values

This method will produce unbiased estimates for the initial margin, but requires sufficiently large n and m values to ensure relatively low variance. McWalter *et al.* (2018) suggest using 1 000 outer simulations and 100 000 inner simulations. This is reasonably implemented when there is a single state variable, but can become cumbersome once X_t becomes multivariate. Due to time and computational limitations, this technique was implemented using 1 000 outer simulations and 10 000 inner simulations. This achieved a balance between time and accuracy of estimates. To improve the effectiveness of the Monte Carlo, a uniform stratification was used, as was done by McWalter *et al.* (2018).

Gaussian Least Squares Monte Carlo

The GLSMC technique, as discussed by McWalter *et al.* (2018), was originally introduced by Anfuso, Aziz, Giltinan and Loukopoulos (2017). This technique finds a functional form for the quantile of the change in portfolio value for a given portfolio value or state variable value. GLSMC can be implemented directly using the state variable when there is a single state variable, but, in the case of a multivariate X_t, it is easier to use the portfolio value as the regressor.

Anfuso *et al.* (2017) assume that the changes in portfolio value over the margin period of risk are normal with a mean of zero. One then uses least squares Monte Carlo to find an expectation for the change in portfolio value squared given a portfolio value. Under the normality assumption, one can use the inverse cumulative distribution function to find the relevant quantile.

As per McWalter *et al.* (2018), the least squares Monte Carlo solves for the measurable function M_2, defined as:

$$M_2(v) = \mathbb{E}^{\mathbb{P}}[(\Delta V_t)^2 | V_t = v],$$

where $\mathbb{P}$ is the real-world measure, V_t is the value of the portfolio at time t, and ΔV_t is the change in the portfolio value over the margin period of risk. Once one has solved for M_2, the quantile for portfolio value changes over the MPOR given the value of portfolio, v, at time t can be solved using the normality assumption of Anfuso *et al.* (2017):

$$Q(\Delta V_t | V_t = v) = \Phi^{-1}(1-\alpha)\sqrt{M_2(v)},$$

where $1-\alpha$ is the confidence level. This is the amount that should be posted as IM at time t, given the portfolio value is equal to v at time t.

The GLSMC technique has the benefit over the nested Monte Carlo technique of being less computationally expensive. McWalter *et al.* (2018) simulated 200 000 values for V_t, and for each of these a single value of $V_{t+\text{MPOR}}$ was simulated. However, the normality assumption does not always hold in practice, and can result in biased estimates for IM.

3.2 Potential Future Exposure

In order to evaluate the effects of initial margin, PFE will be calculated under various collateralisation schemes: posting of variation margin only (termed the 'naive' technique); posting of variation margin and initial margin equal to the closed-form quantile (in the cases where such a form exists); posting of variation margin and initial margin as suggested by nested Monte Carlo; and posting of variation margin and initial margin as suggested by GLSMC.

In all cases, 100 000 paths will be simulated. As required by Basel III, PFE is calculated at the 99% confidence level, which is the same quantile used when calculating initial margin. Basel III also suggests a margin period of risk of 10 business days. Anfuso *et al.* (2017) note that this MPOR may differ in the case of capital exposure calculations.

When computing initial margining, Anfuso *et al.* (2017) suggest using the same paths generated for exposure calculations, as this reduces computational strain. However, in implementing nested Monte Carlo, this was not possible. The exposure path size of 100 000 was too large for an outer loop size, and a smaller exposure path size (like the 1 000 used for the nested Monte Carlo) was too small. As a result, a basic linear interpolation scheme was implemented.

Using the algorithm described in Section 3.1.1 resulted in the following:

Simulated state variables	X_t^1	X_t^2	...	X_t^{1000}
Nested Monte Carlo initial margins	IM_t^1	IM_t^2	...	IM_t^{1000}

In the linear interpolation scheme, these were treated as the known data points. The exposure paths led to a set of 100 000 X_t values, whose relevant initial margin values were treated as the unknowns in the linear interpolation. The standard 1-D data interpolation method, 'inter1', from Matlab was used to perform this.

A similar system was employed for the GLSMC technique. Due to the vastly improved computational speed of the GLSMC algorithm, it was feasible to simulate a larger number of state variable values (in this case, 200 000). However, interpola-

tion and least squares Monte Carlo were performed on the portfolio values, rather than the state variables.

3.3 Margin Value Adjustment

As will be shown, a combination of initial margin and variation margin can prove effective for reducing PFE, and thus the credit risk associated with a contract. However, the posting of these margins leads to an additional cost, since the interest earned on collateral in the margining account tends to be lower than what could be earned with even low-risk investments.

This issue is particularly prevalent in the case of initial margin, since, in 99% of cases, the initial margin posted will exceed the real change in the value of a portfolio. As a result, a margin value adjustment can be calculated, which accounts for the expected future cost of posting initial margin.

McWalter *et al.* (2018) provide the following formula for calculating MVA

$$\text{MVA}_t = B_t \int_t^T \mathbb{E}^{\mathbb{P}}\left[f_s \frac{\text{IM}_s}{B_s} | \mathcal{F}_t \right] ds, \tag{3.1}$$

where MVA_t is the margin value adjustment at time t, B_t is the value of the cash account at time t, T is the maturity time of the contingent claim, f_s is the funding spread at time s, IM_s is the initial margin at time s, and $\mathcal{F}_t$ is the filtration at time t. The funding spread is the difference in interest earned on collateral in the margining account and the rate earned internal to the bank. The difficulty in implementing (3.1) comes from IM_s being an expectation itself. This presents computational difficulties, leading to the need for efficient calculation techniques.

3.4 Sensitivity Techniques

In addition to the nested Monte Carlo and GLSMC techniques, a sensitivity-based technique for calculating initial margin was considered in the simplest case of GBM modelling. This technique involves finding the delta and gamma of a contingent claim, and using these to estimate the 99^{th} percentile of future possible values.

Suppose that δ_t and γ_t are the delta and gamma, respectively, of the claim at time t. Then, assuming that changes in the underlying over the margin period of risk are lognormally distributed, the 99^{th} percentile of portfolio value changes can be estimated using either

$$\text{IM}_t = \delta_t \times (\Delta_{99} S_t), \tag{3.2}$$

or

$$\text{IM}_t = \delta_t \times (\Delta_{99} S_t) + \frac{1}{2} \times \gamma_t \times (\Delta_{99} S_t)^2, \tag{3.3}$$

where $\Delta_{99} S_t$ is the 99^{th} percentile of changes in the underlying, which can be calculated using the lognormal assumption. The estimate given by (3.2) is termed the delta technique, while the one given by (3.3) is termed the gamma technique (this technique also makes use of the delta of claim).

This technique is inspired by the SIMM methodology suggested by ISDA (International Swap and Derivatives Association, 2018), which details the calculation and use of sensitivities for finding IM values. The estimation of these greek values is a process in and of itself, and so we will only consider them in a simple model, addressed in the next chapter.

Chapter 4

Call Option Results

The techniques and procedures detailed in Chapter 3 were implemented in two cases: an FX call option where the exchange rate is modelled using geometric Brownian motion; and an interest rate swap (IRS) where the short rate is modelled using a two-factor model. This chapter considers the call option, which is investigated using an unsophisticated model for illustrative purposes. The next chapter focuses on a model and instrument that are more likely to be implemented in practice.

The FX call option being modelled is presented as a call on USD (the foreign currency) using ZAR (the domestic currency). The payoff of this option at maturity, time T, is given by:

$$\max\{N \times (\mathrm{E}_T - K), 0\},$$

where N is the nominal number of USD, E_T is the cost of one USD in terms of ZAR at maturity, and K is the strike price of the option.

The parameters and assumptions for the contingent claim being considered here mirror those shown in Chapter 1, namely: exchange rates follow standard geometric Brownian motion (GBM); notional of one million USD with maturity of one year; spot price for USD of 13 ZAR; risk-free continuously compounded interest rates, with r_{ZAR} of 8%, and r_{USD} of 1.5%; volatility of stochastic process, σ, of 30%.

We assume no trading costs or tax, and fractional and negative holdings permissible. Under GBM, one can then simulate the exchange rate at time t from time s using

$$\mathrm{E}_t = \mathrm{E}_s \times \exp\left[\left(r_{ZAR} - r_{USD} - \frac{\sigma^2}{2}\right) \times (t-s) + \sigma\sqrt{t-s}Z\right], \tag{4.1}$$

where E_s is the exchange rate at time s, and Z is a standard normal random variable. At this point, it is necessary to note that the exchange rate has been simulated under the risk-neutral measure — the same measure under which pricing is done. The drift in (4.1) has its form to ensure that converting money and then

investing it in the foreign cash account gives a martingale when domestically discounted. Some bias can emerge as a result of using the risk-neutral measure, and, strictly speaking, one should simulate the paths used for PFE calculations under the real-world measure — this is because PFE refers to quantiles in the real-world. However, this simplification is acceptable, since the risk-neutral distribution still captures a spread of possibilities, and is a sufficient proxy, especially over a short margin period of risk.

Given the assumptions above, it is possible to find the price of the FX call option at time t, random variable C_t, given the simulated exchange rate at time t using the Black-Scholes formula as

$$C_t = \Phi(d_1)\mathrm{E}_t - \Phi(d_2)Ke^{-r(T-t)}, \tag{4.2}$$

where K is the strike price, $r = r_{ZAR} - r_{USD}$, T is the maturity, Φ is the standard normal cumulative distribution, and

$$\begin{aligned} d_1 &= \frac{1}{\sigma\sqrt{T-t}}\Big[\ln\Big(\frac{\mathrm{E}_t}{K}\Big) + \Big(r + \frac{\sigma^2}{2}\Big)(T-t)\Big] \\ d_2 &= d_1 - \sigma\sqrt{T-t}. \end{aligned} \tag{4.3}$$

With this pricing formula, one can simulate E_t values forward by the margin period of risk (MPOR), giving values for $E_{t+\text{MPOR}}$ and thus $C_{t+\text{MPOR}}$. This is what can be used to perform a nested Monte Carlo simulation or the GLSMC technique to find estimates for the initial margin.

For this model, the delta of the option, δ_t, is given simply by $\Phi(d_1)$, while the gamma, γ_t, is given by

$$\frac{\phi(d_1)}{E_t\sigma\sqrt{T-t}}, \tag{4.4}$$

where ϕ is the standard normal density function.

When using GBM to model exchange rates, there is an analytical solution for the initial margin required at time t given the value of the underlying at time t. Noting the formula given by (4.1), the 99^{th} percentile for the exchange rate after the margin period of risk is given by

$$\mathrm{E}_{\text{percentile}} = \mathrm{E}_t \times \exp\Big[\Big(r_{ZAR} - r_{USD} - \frac{\sigma^2}{2}\Big) \times \text{MPOR} + \sigma\sqrt{\text{MPOR}}Z_{99}\Big], \tag{4.5}$$

where Z_{99} is the 99^{th} percentile for the standard normal distribution (i.e. $0.99 = \Phi(Z_{99})$), and $\mathrm{E}_{\text{percentile}}$ is the 99^{th} percentile for $\mathrm{E}_{t+\text{MPOR}}$ given E_t. This is because the future exchange rate is monotonic increasing with respect to Z. Now, noting that the value of an FX call option is monotonic with respect to the underlying exchange rate, the 99^{th} percentile for the call option value following the MPOR given the

underlying at time t can be calculated using (4.2), where one uses $\mathrm{E}_{\text{percentile}}$ as the exchange rate. Note that if $T - t <$ MPOR, one would consider percentiles of underlying values at maturity, and not at $t +$ MPOR. In this case, (4.5) would use $T-t$ as opposed to MPOR for the elapsed time, and the call option would be valued at its exercise value (since one would be at maturity).

Figure 4.1 shows initial margin values for various possible FX option values halfway through the lifespan of the option (6 months). In this example, the option is in-the-money at inception with a strike of 11.5 ZAR per USD, and values are shown on a single claim, i.e., on a notional of 1 USD. The black points show the analytical IM values — these represent the exact 99% level for the change in option value over the margin period of risk. The blue circles show the estimated IM values using a nested Monte Carlo algorithm. The red crosses show the estimated IM values using the GLSMC algorithm.

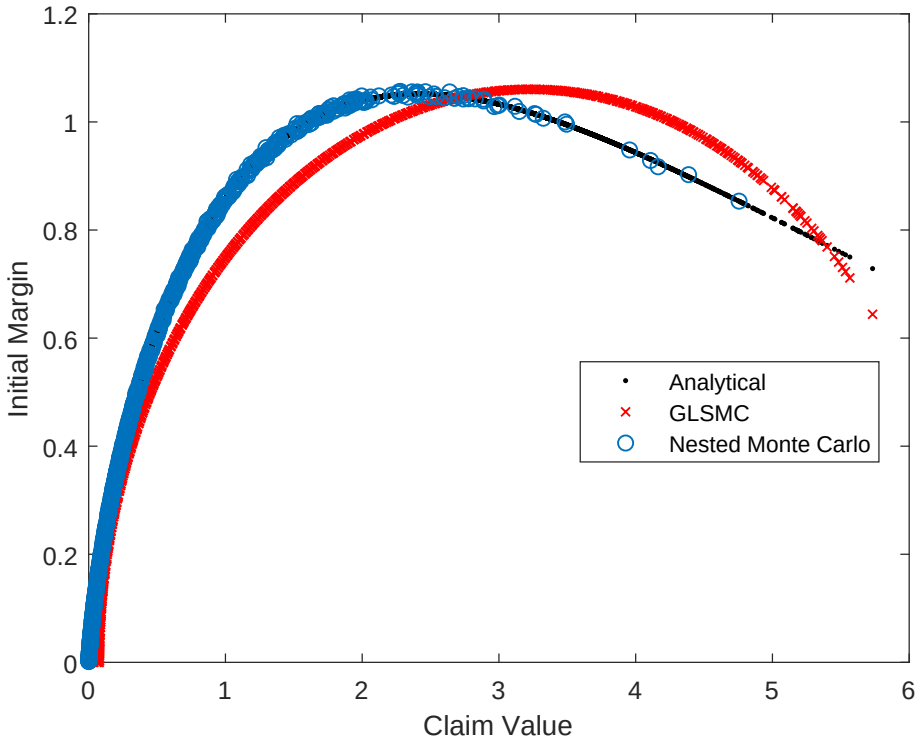

Fig. 4.1: Initial margin values for various contingent claim values (6 months).

This was implemented similarly as McWalter *et al.* (2018). An outer sample size for the nested Monte Carlo of 1 000 was used, with an inner sample size of 200 000, with uniform stratification. For the GLSMC and analytical cases, sample sizes of 200 000 were used. Given the fact that the nested Monte Carlo provides

an unbiased estimate for the IM, it is unsurprising that there is little discrepancy between its estimates and the closed-form IM values. It is worth noting that the outer sample size of 1 000 did not result in extremely large contingent claim values.

The GLSMC provides a less accurate estimate for IM values, reflective of the limitations of the normality assumption it hinges on. Over the entire sample size, it exhibits an error of roughly 35%. However, one can observe that the error is smaller in the tails — this is most relevant when considering PFE calculations, since one would usually be working within tails of possible portfolio values.

Figure 4.2 shows the PFE of an FX call option that is in-the-money at inception. The initial spot price for USD is 13 ZAR; the FX call option shown in this figure has a strike of 11.5 ZAR per USD. The light blue line shows the PFE of the option over its lifespan when only variation margin is posted — the 'naive' technique. The orange line with squares shows the PFE when both variation and initial margin are posted, with the initial margin estimate using nested Monte Carlo. The yellow line with stars shows variation and initial margin postings, with initial margin estimated using the GLSMC technique. Finally, the purple with circles and green dashed lines, respectively, show the PFE when initial margin is calculated using the delta and gamma techniques, as explained in Chapter 3, and delta values calculated using $\Phi(d_1)$ and gamma values using (4.4).

In all cases, the PFE of the option is the same over an initial period. The length of this period is equal to the margin period of risk, as any collateral posted at inception has yet to be accounted for. Following this, collateral begins to have an effect. As evidenced by the naive technique, variation margin itself is not sufficient to eliminate PFE. Regardless of which technique is used to calculate IM, there is a large reduction in the level of PFE at any time following the initial MPOR. The gamma technique is completely effective in eliminating PFE. The nested Monte Carlo technique is the next most effective, likely as a result of it being an unbiased estimation technique. Due to limited sample size, it performs marginally worse than the second order approximation given by the gamma technique. The error present in the GLSMC technique illustrated in Figure 4.1 is apparent in its lower effectiveness in reducing PFE. The delta technique provides a better estimate than the GLSMC technique, but is inferior to both the nested Monte Carlo and gamma technique. Both the delta and GLSMC techniques have relatively constant absolute levels of bias over the lifespan of the option. While not graphed, closed-form initial margin values were calculated as per (4.5). These closed-form values fully reduced PFE, as one would expect.

Figure 4.3 shows the percentage error in IM values using the various techniques. The IM estimates were compared to the closed-form IM values, and were measured

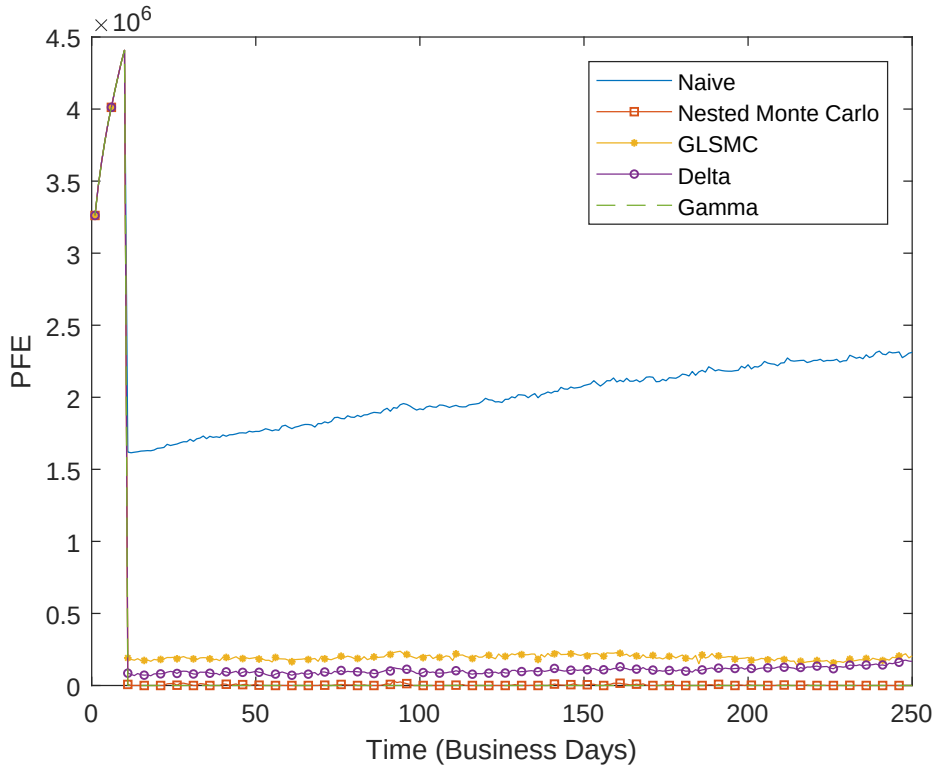

Fig. 4.2: Potential future exposure for in-the-money FX call option.

in terms of underestimating IM.

In the case of the nested Monte Carlo, GLSMC, and delta techniques, their IM estimates were always less than the closed-form IM values. The gamma technique overestimated IM values near inception, and underestimated them towards maturity. The nested Monte Carlo, delta, and gamma techniques have relatively smooth error plots, while the GLSMC grows less smooth over time. The erratic behaviour of the GLSMC towards maturity is likely due to the normality assumption becoming less applicable as the time to maturity in the option is reduced. None of the techniques exhibited an absolute percentage error above 30%, with the nested Monte Carlo and gamma techniques remaining within 10% for most of the lifespan of the option.

Table 4.1 summarises the key aspects of the techniques tested. The first column shows the maximum PFE per unit notional over the lifespan of the option (excluding the initial period before collateral has come into effect); the second column shows the average PFE per unit notional over the lifespan of the option (again, excluding the beginning period); the final column shows the average IM error over the lifespan of the option when compared to the closed-form IM.

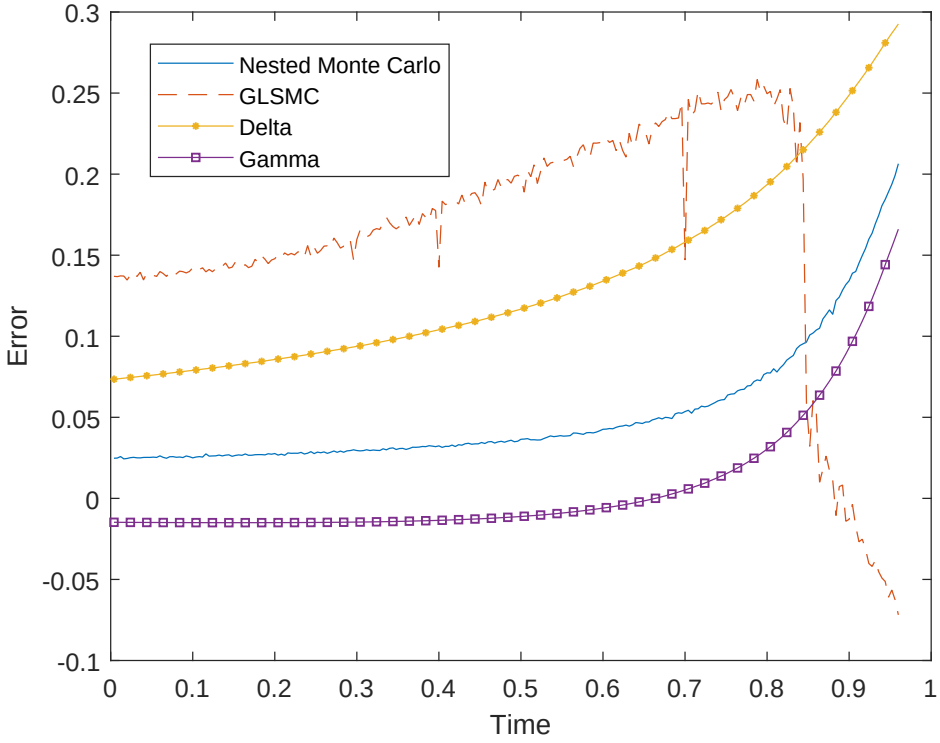

Fig. 4.3: Percentage error in initial margin for in-the-money FX call option.

Tab. 4.1: In-the-money PFE results.

Technique	Max PFE	Average PFE	IM Error
Naive	2.3177	2.0010	NA
Nested Monte Carlo	0.0289	0.0033	5.1509%
GLSMC	0.2351	0.1939	16.2681%
Delta	0.1692	0.1049	13.3334%
Gamma	0	0	0.6943%

As is apparent from Figure 4.2, all IM techniques are effective in reducing PFE, to varying degrees. The gamma technique resulted in zero PFE; however, as can be seen in Figure 4.3, it sometimes overestimated IM values. The nested Monte Carlo technique was next most effective, followed by the delta and GLSMC techniques — all consistent with Figures 4.2 and 4.3.

Using the simulated exchange paths and their relevant initial margin estimates,

the margin value adjustment was calculated. In the formula

$$\text{MVA}_t = B_t \int_t^T \mathbb{E}^{\mathbb{P}}\left[f_s \frac{\text{IM}_s}{B_s} | \mathcal{F}_t \right] ds,$$

the funding spread, f_s, was taken as 2%. The expectation in the integral was estimated using the mean of the estimated IM values, after which the integral was estimated using a trapezoidal technique.

Table 4.2 shows the results of the MVA calculations. The first column lists the technique considered; the second column shows the MVA per unit notional calculated; the third shows the percentage error present in the technique when compared to the MVA using closed-form IM values. This percentage error is considered in terms of underestimating the MVA.

Tab. 4.2: In-the-money MVA estimates.

Technique	MVA	MVA Error
Closed-form	0.03144	0
Nested Monte Carlo	0.03141	0.09777%
GLSMC	0.02811	10.57948%
Delta	0.02914	7.32266%
Gamma	0.03191	-1.50712%

The closed-form IM values give an MVA per unit notional of 0.03144. The nested Monte Carlo, GLSMC, and delta IM-estimation techniques all result in lower estimates for MVA, while the gamma technique overestimates it. The nested Monte Carlo provides the best estimate, being 9.78 basis points too low. The GLSMC and delta techniques are roughly 10.6 and 7.3 percent too low, respectively, while the gamma is roughly 1.5 percent too large.

The MVA is an additional cost that should be considered when entering into a contract. The writer of this option should increase the premium by 0.03144 per unit notional, to account for the cost of implementing an initial margin.

While the nested Monte Carlo provided the best MVA estimate, and generally performed well in other aspects, it was the most computationally expensive technique. It had a runtime of 552 seconds, many times larger than the other techniques. The GLSMC was next slowest, requiring 72 seconds to complete. The delta and gamma techniques ran for approximately 3 and 6 seconds respectively. Overall, the gamma technique provides the best balance between timing and accuracy. However, the gamma and delta techniques would likely prove far more difficult to implement in more complex scenarios, such as portfolios including numerous

correlated underlyings. They thus represent idealisations — their calculation and implementation could prove far more cumbersome in practice.

To conclude the chapter, we briefly consider the case where the option is written out-of-the-money instead, to illustrate the consistency of the various techniques. Figure 4.4 shows the PFE of the same option above struck at 16 ZAR per USD.

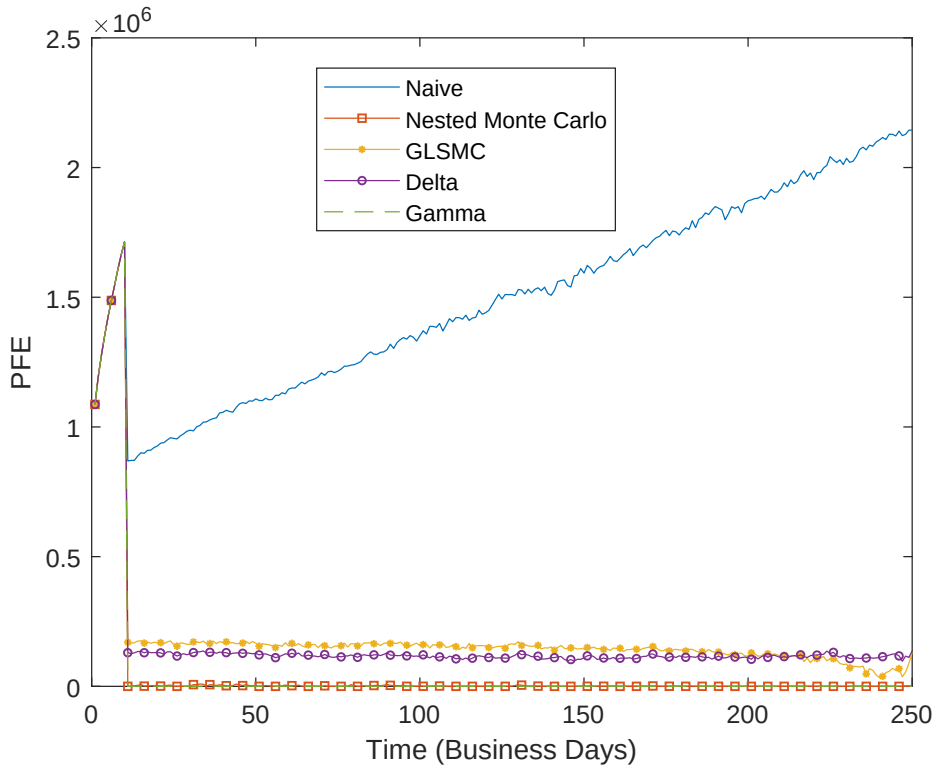

Fig. 4.4: Potential future exposure for out-of-the-money FX call option.

In this case, the naive technique shows less effectiveness over the lifespan of the option. The PFE exhibits a sharp upward trend after its initial drop following the first MPOR. While this trend is present in the in-the-money option, it is more pronounced in the out-of-the-money case. As previously, all techniques for calculating initial margin prove effective, to varying degrees. The bias of the GLSMC and delta techniques are similar to that seen for the in-the-money option. The nested Monte Carlo technique has only slight bias, as a result of a limited sample size. Once again, the gamma technique is completely effective in reducing PFE.

Figure 4.5 replicates Figure 4.3 for the out-of-the-money option.

The results are similar to the in-the-money option, with the progressively greater error towards maturity and the GLSMC showing erratic behaviour towards the end. Overall, the level of error is greater on average. This can also be seen in Ta-

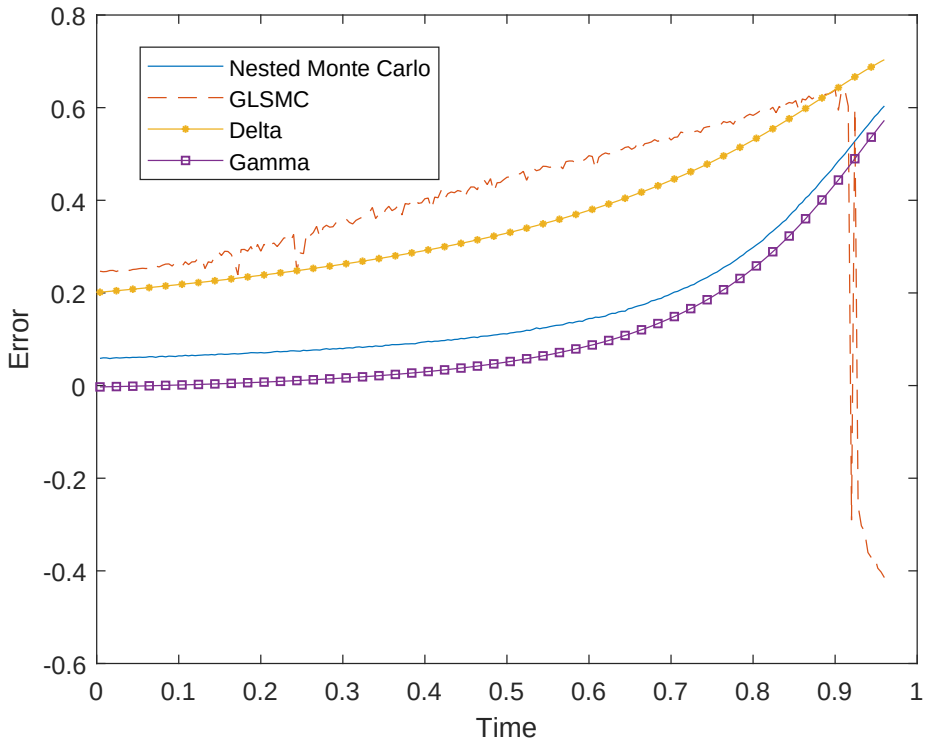

Fig. 4.5: Percentage error in initial margin for out-of-the-money FX call option.

ble 4.3, which replicates Table 4.1.

The IM errors tended to be greater than those seen for the in-the-money option. The last illustration of this is in Table 4.4. Once again, the gamma technique is the only one to overestimate MVA. Nested Monte Carlo still provides the closest estimate, but all estimates show greater error than for the in-the-money option. The overall MVA is smaller for the out-of-the-money option, but would still be necessary to consider as the writer of such an option.

Tab. 4.3: Out-of-the-money PFE results.

Technique	Max PFE	Average PFE	IM Error
Naive	2.1496	1.5015	NA
Nested Monte Carlo	0.0156	0.0009	17.1795%
GLSMC	0.1823	0.1388	39.2067%
Delta	0.1447	0.1150	36.6882%
Gamma	0	0	11.5319%

Tab. 4.4: Out-of-the-money MVA estimates.

Technique	MVA	MVA Error
Closed-form	0.01679	0
Nested Monte Carlo	0.01674	0.26344%
GLSMC	0.01383	17.63246%
Delta	0.01391	17.17167%
Gamma	0.01746	-4.00879%

Chapter 5

Interest Rate Swap Results

This chapter considers the implementation of an initial margin in the case of an interest rate swap (IRS). Due to the more complicated nature of this type of contract, only the nested Monte Carlo and GLSMC techniques were applied.

5.1 Interest Rate Swap Valuation

The interest rate swap contingent claim being modelled is a fixed-for-floating IRS, where payments termed as fixed are known from inception and payments termed as floating are only known one period prior to their being paid. One party makes the floating payments to the other, and in return receives the fixed payments. The IRS is defined by a set of reset and payment dates, $\{T_0 = 0, T_1, T_2, ..., T_n = T\}$, a fixed rate, K, and nominal amount, N. The accrual period between consecutive T_i values is usually constant, i.e. $T_{i+1} - T_i = T_{j+1} - T_j = \Delta T$ for all (i, j), where ΔT is expressed as a year fraction. On each of the reset dates, $T_0, T_1, ..., T_{n-1}$, the floating interest which is to be paid or received at the next payment date is set to the market-observed rate over the coming period. This is usually an IBOR rate, such as LIBOR or JIBOR.

Suppose $L(T, T + \Delta T)$ represents the market-observed interest rate at time T, which applies from time T to $T + \Delta T$. Then, the floating payment of the IRS will be

$$NL(T, T + \Delta T)\Delta T,$$

at time $T + \Delta T$. By a no-arbitrage argument, the present value of this payment at time t is given by

$$NL(t, T, T + \Delta T)\Delta T B(t, T + \Delta T), \tag{5.1}$$

where $L(t, T, T + \Delta T)$ is the time t forward rate applying from time T to time $T + \Delta T$, and $B(t, T + \Delta T)$ is the discount factor applying at time t for time $T + \Delta T$. Again, through a no-arbitrage argument, the value of $L(t, T, T + \Delta T)$ can be found

from time t discount factors to be

$$L(t,T,T+\Delta T)=\frac{1}{\Delta T}\Big(\frac{B(t,T)}{B(t,T+\Delta T)}-1\Big). \tag{5.2}$$

Using (5.1), the value of the floating leg of an IRS at time $T_i \leq t < T_{i+1}$ is

$$N\Delta T\left(L(T_i,T_{i+1})B(t,T_{i+1})+\sum_{j=i+1}^{n-1}L(t,T_j,T_{j+1})B(t,T_{j+1})\right). \tag{5.3}$$

Note that the payment due on the next payment date is known in advance.

All fixed payments made are the same, and known from inception. Suppose the fixed payments are made on some subset of $\{T_1,T_2,...,T_n\}$, $\{T_{f_1},T_{f_2},...,T_{f_m}\}$, with $T_{f_{i+1}}-T_{f_i}=\Delta T_f$ for all i. Then, the fixed payment made on each T_{f_i} is

$$NK\Delta T_f,$$

where K is the fixed rate. Thus, the value of the fixed payments at $T_{f_i} \leq t < T_{f_{i+1}}$ is

$$NK\Delta T_f\sum_{j=i+1}^{m}B(t,T_{f_j}). \tag{5.4}$$

In many cases, the fixed rate, K, will be set such that the value of the IRS is zero at inception. The main IRS considered here will have fixed and floating payments on the same dates; however, an IRS where this is not the case will be briefly analysed for illustrative purposes.

5.2 G2++ Model

In order to analyse an IRS, an interest rate model is needed. The model considered is the G2++ model described by Brigo and Mercurio (2006), with derivations courtesy of Lopes (2018). This is a more general model than, for example, the Hull-White model, and can model path-dependent claims more realistically. With this model, the instantaneous short rate under the risk-neutral measure, $\mathbb{Q}$, at time t is given by

$$r(t)=x_1(t)+x_2(t)+\varphi(t), \quad r(0)=r_0, \tag{5.5}$$

where $x_1(t)$ and $x_2(t)$ are driven by the stochastic differential equations (SDEs)

$$\begin{aligned} dx_1(t)&=-\kappa_1x_1(t)dt+\sigma_1dW_1(t), \quad x_1(0)=0,\\ dx_2(t)&=-\kappa_2x_2(t)dt+\sigma_2dW_2(t), \quad x_2(0)=0, \end{aligned}$$

with $W_1(t)$ and $W_2(t)$ being $\mathbb{Q}$-Brownian motions with a correlation of ρ. As detailed by Lopes (2018), the function φ is defined on a time interval $[0, T^*]$, with $\varphi(0) = r_0$ and

$$\begin{aligned}\varphi(T) = f(0,T) + \frac{\sigma_1^2}{2\kappa_1^2}(1 - e^{-\kappa_1 T})^2 + \frac{\sigma_2^2}{2\kappa_2^2}(1 - e^{-\kappa_2 T})^2 \\ + \rho\frac{\sigma_1\sigma_2}{\kappa_1\kappa_2}(1 - e^{-\kappa_1 T})(1 - e^{-\kappa_2 T}),\end{aligned}$$

where $T \in [0, T^*]$ and $f(0,T)$ is the forward rate applicable from time 0 to T. This ensures that the model described by (5.5) fits the initial term structure of interest rates.

The random variable $Y(t,T) = \int_t^T \big(x_1(u) + x_2(u)\big)du$ is distributed normally (Lopes, 2018) with a mean of

$$\begin{aligned}\mathbb{E}[Y(t_1,t_2)|\mathcal{F}_{t_1}] &= \frac{1 - e^{-\kappa_1(t_2-t_1)}}{\kappa_1}x_1(t_1) + \frac{1 - e^{-\kappa_2(t_2-t_1)}}{\kappa_2}x_2(t_1) \\ &=: M(t_1,t_2),\end{aligned}$$

and a variance of

$$\begin{aligned}\mathbb{V}[Y(t_1,t_2)|\mathcal{F}_{t_1}] = \frac{\sigma_1^2}{\kappa_1^2}\Big[(t_2 - t_1) - A_1(t_1,t_2) - \frac{\kappa_1}{2}A_1^2(t_1,t_2)\Big] \\ + \frac{\sigma_2^2}{\kappa_2^2}\Big[(t_2 - t_1) - A_2(t_1,t_2) - \frac{\kappa_2}{2}A_2^2(t_1,t_2)\Big] \\ + 2\rho\frac{\sigma_1\sigma_2}{\kappa_1\kappa_2}\Big[(t_2 - t_1) - A_1(t_1,t_2) - A_2(t_1,t_2) \\ - \frac{e^{-(\kappa_1+\kappa_2)(t_2-t_1)} - 1}{\kappa_1 + \kappa_2}\Big] \\ =: V(t_1,t_2),\end{aligned}$$

where

$$A_i(t_1,t_2) = \frac{1}{\kappa_i}\Big(1 - e^{-\kappa_i(t_2-t_1)}\Big).$$

This allows one both to simulate the numeraire as well as to price zero-coupon bonds (ZCB). As per Lopes (2018), the price of a ZCB at time t maturing at time T is given by

$$P(t,T) = \exp\left(-\int_t^T \varphi(u)du - M(t,T) + \frac{1}{2}V(t,T)\right).$$

Using these ZCB prices as discount factors, one is able to value interest rate swaps at any time t using (5.2), (5.3), and (5.4).

Since $M(t_1,t_2)$ is given in terms of $x_1(t_1)$ and $x_2(t_1)$, it is necessary to simulate paths jointly for $x_1(t)$ and $x_2(t)$ to find a path realisation for the IRS. Additionally,

in order to find estimates for MVA, one requires a jointly simulated realisation for $Y(0,t)$ to find the value of the numeraire for a given short rate path. Lopes (2018) derives the covariance structure of these three variables as follows

$$\begin{aligned}\mathbb{C}_{t_1}[x_1(t_2), Y(t_1,t_2)] = \rho\frac{\sigma_1\sigma_2}{\kappa_1+\kappa_2}\left[A_2(t_1,t_2) + \frac{1}{\kappa_2}\left(e^{-(\kappa_1+\kappa_2)(t_2-t_1)} - e^{-\kappa_1(t_2-t_1)}\right)\right] \\ + \frac{\sigma_1^2}{2\kappa_1}\left[A_1(t_1,t_2) + \frac{1}{\kappa_1}\left(e^{-2\kappa_1(t_2-t_1)} - e^{-\kappa_1(t_2-t_1)}\right)\right],\end{aligned}$$

with $\mathbb{C}_{t_1}[x_2(t_2), Y(t_1,t_2)]$ being of the same form, and

$$\begin{aligned}\mathbb{C}_{t_1}[x_1(t_2), x_2(t_2)] = \rho\frac{\sigma_1\sigma_2}{\kappa_1\kappa_2}\Bigg[(t_2-t_1) - A_1(t_1,t_2) - A_2(t_1,t_2) \\ - \frac{e^{-(\kappa_1+\kappa_2)(t_2-t_1)} - 1}{\kappa_1+\kappa_2}\Bigg].\end{aligned}$$

Now, noting that

$$\mathbb{V}[x_1(t_2)|\mathcal{F}_{t_1}] = \frac{\sigma_1^2}{\kappa_1^2}\left[(t_2-t_1) - A_1(t_1,t_2) - \frac{\kappa_1}{2}A_1^2(t_1,t_2)\right],$$

and that $\mathbb{V}[x_2(t_2)|\mathcal{F}_{t_1}]$ has the same form, the 3×3 t_2 cross-correlation matrix for $(x_1(t_2), x_2(t_2), Y(t_1,t_2))$ at time t_1 is given by

$$\begin{bmatrix} 1 & \frac{\mathbb{C}_{t_1}[x_1(t_2),x_2(t_2)]}{\sqrt{\mathbb{V}[x_1(t_2)|\mathcal{F}_{t_1}]\mathbb{V}[x_2(t_2)|\mathcal{F}_{t_1}]}} & \frac{\mathbb{C}_{t_1}[x_1(t_2),Y(t_1,t_2)]}{\sqrt{\mathbb{V}[x_1(t_2)|\mathcal{F}_{t_1}]\mathbb{V}[Y(t_1,t_2)|\mathcal{F}_{t_1}]}} \\ & 1 & \frac{\mathbb{C}_{t_1}[x_2(t_2),Y(t_1,t_2)]}{\sqrt{\mathbb{V}[x_2(t_2)|\mathcal{F}_{t_1}]\mathbb{V}[Y(t_1,t_2)|\mathcal{F}_{t_1}]}} \\ & & 1 \end{bmatrix},$$

which is presented in the same form as Lopes (2018) by omitting the lower entries in the symmetric matrix.

One can perform a Cholesky decomposition on this matrix to find a lower triangular matrix L, such that, if Z is a $3\times n$ matrix of standard normal random variables, $X = LZ$ will be a $3\times n$ matrix of correlated standard normal random variables. Then, using the fact that $x_1(t_2)$, $x_2(t_2)$, and $Y(t_1,t_2)$ are all normally distributed given $\mathcal{F}_{t_1}$, they can be jointly simulated n times using

$$\begin{aligned} x_1(t_2) &= \mathbb{E}[x_1(t_2)|\mathcal{F}_{t_1}] + \sqrt{\mathbb{V}[x_1(t_2)|\mathcal{F}_{t_1}]}X(1) \\ x_2(t_2) &= \mathbb{E}[x_2(t_2)|\mathcal{F}_{t_1}] + \sqrt{\mathbb{V}[x_2(t_2)|\mathcal{F}_{t_1}]}X(2) \\ Y(t_1,t_2) &= \mathbb{E}[Y(t_1,t_2)|\mathcal{F}_{t_1}] + \sqrt{\mathbb{V}[Y(t_1,t_2)|\mathcal{F}_{t_1}]}X(3), \end{aligned}$$

where $X(i)$ is the i^{th} row of X, and

$$\mathbb{E}[x_i(t_2)|\mathcal{F}_{t_1}] = e^{-\kappa_i(t_2-t_1)}x_i(t_1).$$

Lopes (2018) presents this in a slightly different way, as simulating one set of x_1, x_2, and Y at n time points, but the derivations follow in the same way.

In implementing this model, the parameters chosen were equivalent to those in Lopes (2018), namely: $\kappa_1 = 0.1$, $\kappa_2 = 0.25$, $\sigma_1 = 0.025$, $\sigma_2 = 0.02$, and $\rho = -0.6$. The initial term structure of interest rates, to which φ is fitted, is described by a Vasicek (1977) short rate model, with mean reversion rate 0.1, mean reversion level 0.09, volatility 0.04, and initial short rate 0.07. Values for $\varphi(t)$ are found for $t \in \{\frac{1}{24}, \frac{2}{24}, \frac{3}{24}, \ldots, T\}$, where T is the maturity of the IRS considered and t is a year fraction. These values are used as they correspond to both the margining periods, and adjacent times differ by the margin period of risk.

5.3 Exposure Using the G2++ Model

Using the model and results of the previous section, as well as the IRS valuation methodology of Section 5.1, we are able to evaluate potential future exposure and initial margin in this model. At certain points, it is necessary to estimate values for $\int_t^s \varphi(u)du$. In these cases, the integral is estimated using the trapezoidal rule and the time increments described at the end of the previous section.

Firstly, in order to illustrate the functionality of this model, the expected exposure profile of a fixed-for-floating interest rate swap is considered. The swap considered is similar to the one discussed in Andersen, Pykhtin and Sokol (2017b). This swap has floating payments every 3 months, and fixed payments every 6 months, with variation margin posted every 10 business days and a margin period of risk of 10 days. No initial margin is posted, and it is assumed that both parties cease making margin payments once one of them has defaulted — this is termed the classical+ model in Andersen *et al.* (2017b). In all subsequent examples, this is the assumption that is made.

Figure 5.1 shows the expected exposure of the IRS described above with a maturity of 5 years and fixed rate set such that the initial value of the swap is zero. The exposure profile behaviour shows similar characteristics to the first figure in Andersen *et al.* (2017b). There are downward spikes at the dates when only floating payments are made, and larger upward spikes when fixed payments are made. As explained by Andersen *et al.* (2017b), this is expected behaviour, since the fixed payments will generally be larger, as they are made more infrequently.

This example illustrates the adequacy of this model. Henceforth, the IRS considered will have fixed and floating payments on the same dates (quarterly) over a five year period. The fixed rate is set such that the initial value of the IRS is zero. Margin is posted every 10 business days, which is equal to the margin period of

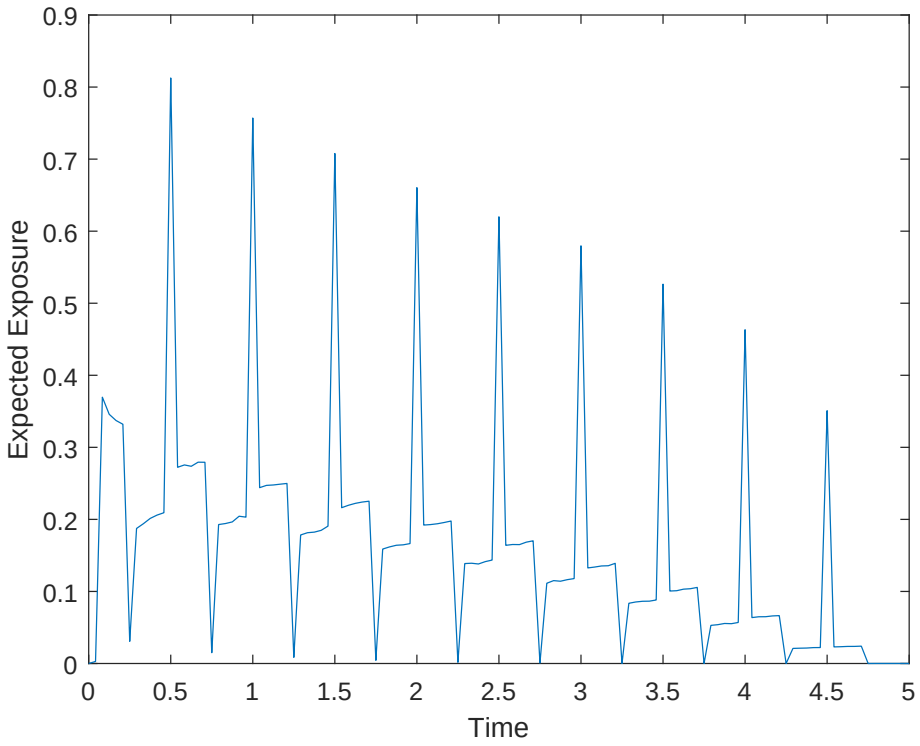

Fig. 5.1: Expected exposure of IRS.

risk, with the contract being valued on the same days. As was the case in Chapter 4, 100 000 simulations were used, and initial margin values and PFE estimated on these. All figures are considered from the viewpoint of the party who receives floating and pays fixed.

Figure 5.2 shows the first 20 of the 100 000 simulated IRS value paths over the lifespan of the claim. Given that this is from the viewpoint of the floating-receiver, positive IRS values indicate points at which the value of floating payments exceed the value of fixed payments. The value is initially at zero for all paths, and the values are each constant over the final period. This is because the final payment is known in advance, since the floating payment is always set one period before it is paid.

As proves to be a common feature of all figures concerning the swap, there are a number of spikes over the lifespan of the claim. These spikes are as a result of payments made during the contract — a feature not seen in the contract considered in Chapter 4.

Figure 5.3 shows the potential future exposure of the IRS over its lifespan. The light blue line is the PFE when no collateral is posted; the orange dashed line is

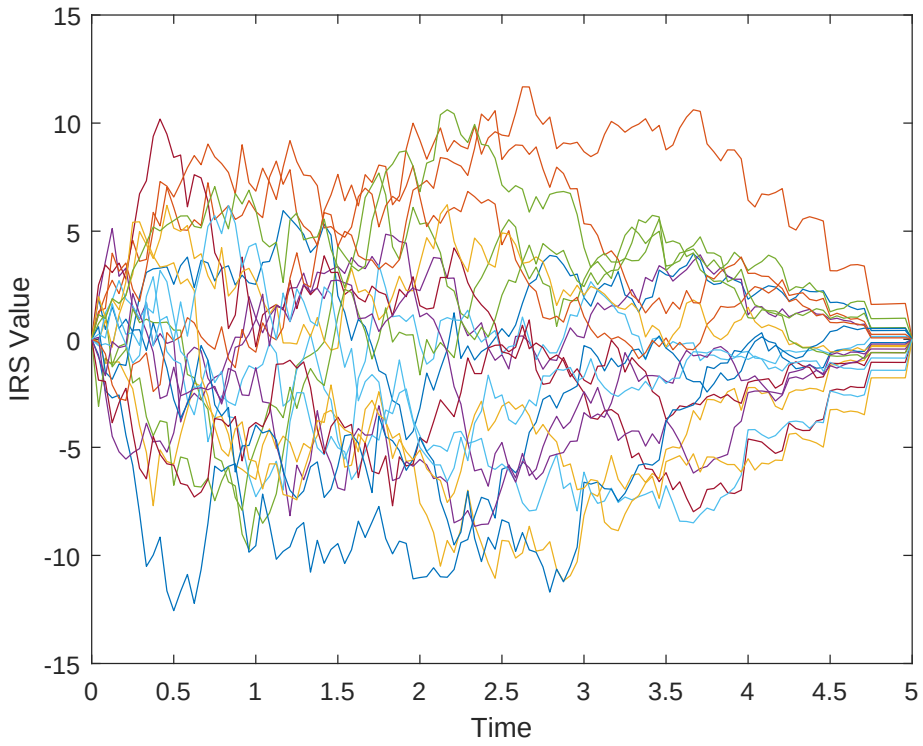

Fig. 5.2: IRS value over time.

when only variation margin is posted; and the yellow with squares and purple with stars lines are when initial margin using nested Monte Carlo and GLSMC, respectively, is posted in addition to variation margin.

In the case of no collateral, there are two opposing forces affecting PFE: as time progresses, there is more potential for the claim to move deeper into the money, making future payments more valuable; however, amortisation reduces the number of future payments, which lowers the value of the claim. By the end of the contract, with only one payment remaining, the value is close to zero (the starting value).

When only variation margin is posted, there is a significant reduction in PFE. In comparison to the FX call option previously considered, here the PFE tends steadily downwards, rather than slowly increasing over time. This is as a result of the amortisation effect, as seen in the no collateral case. Following the second to last payment, the PFE is zero, since, at this time, the final payment would be known, hence the final value of the IRS — due to this, variation margin would always prove sufficient for covering future exposure.

As was seen in the FX call option, using nested Monte Carlo to calculate IM is

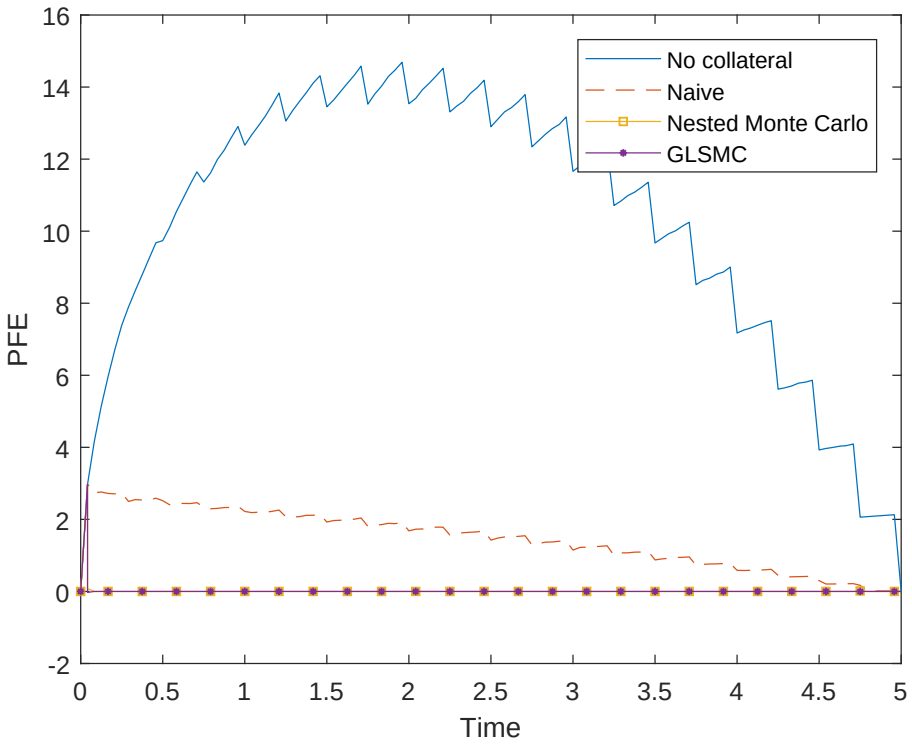

Fig. 5.3: IRS PFE over time.

essentially completely effective in eliminating PFE. While not clearly visible in the figure, the first posted nested Monte Carlo-estimated IM was slightly too small, but this is most likely an artefact of a limited sample size. Unlike in the FX call option, the GLSMC technique for calculating initial margin proves largely effective.

Since the GLSMC algorithm assumes portfolio value changes over the margin period of risk are normal, its effectiveness suggested that changes under the G2++ model are roughly normal. Figure 5.4 show the sample portfolio value changes over the MPOR one quarter of the way through the contract, halfway through the contract, and three quarters through, as one moves down. The figures on the left show histograms of the portfolio value changes, with the red border line indicating a normal distribution fit. The figures on the right show QQ-plots of the sample distributions, where the red dashed lines are the expected number of standard normal quantiles under a normal distribution.

At all three points in time, the data appear close to normal with zero mean. When considering the QQ-plots, there are very few outliers. These together confirm the accuracy of the normality assumption made by the GLSMC technique. Under both the G2++ model and the Vasicek model, which describes the initial

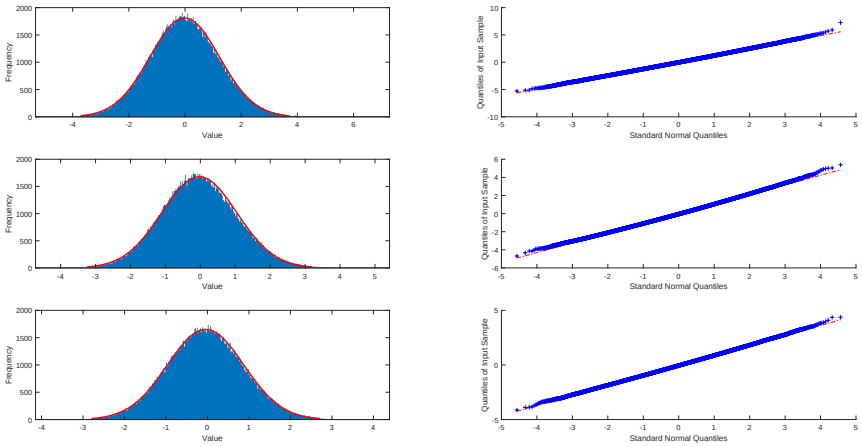

Fig. 5.4: IRS value changes over MPOR.

term structure, the short rate has a normal distribution. This is the likely cause for the changes in IRS values to be normal, since they are driven by the short rate.

Having concluded that the GLSMC assumption is valid, one can consider how the two techniques performed relative to each other. For the contract and model being considered, there is no closed-form solution, leading to direct comparison being the best approximation for accuracy.

Figure 5.5 shows the percentile at which the initial margins for the two techniques are insufficient. The blue line graphs the nested Monte Carlo; the orange dashed line shows the GLSMC; the yellow with stars line shows the 99^{th} percentile — this is the level at which the initial margin is expected to fail.

As can be seen, up until just before the conclusion of the IRS, both techniques are slightly over-estimating IM. Both have spikes at payment times, with the GLSMC spikes being more severe, and upwards as opposed to downwards. However, towards the end of the contract, the GLSMC has a sharp downward spike, where initial margin is insufficient. This is likely due to the normality assumption being less accurate at the extreme end of the contract. It is important to note, however, that these values do not explain the magnitude of over- or underestimation of the IM. For example, at the second to last payment, where both the nested Monte Carlo and GLSMC techniques give too small an initial margin, the PFE is roughly 0.0004 in both cases. This small magnitude is why the spike seen in Figure 5.5 does not lead to a discernible spike in Figure 5.3. In totality, the two techniques perform close to the level at which they are expected.

Figure 5.6 shows the average initial margin for all simulations over time. Once

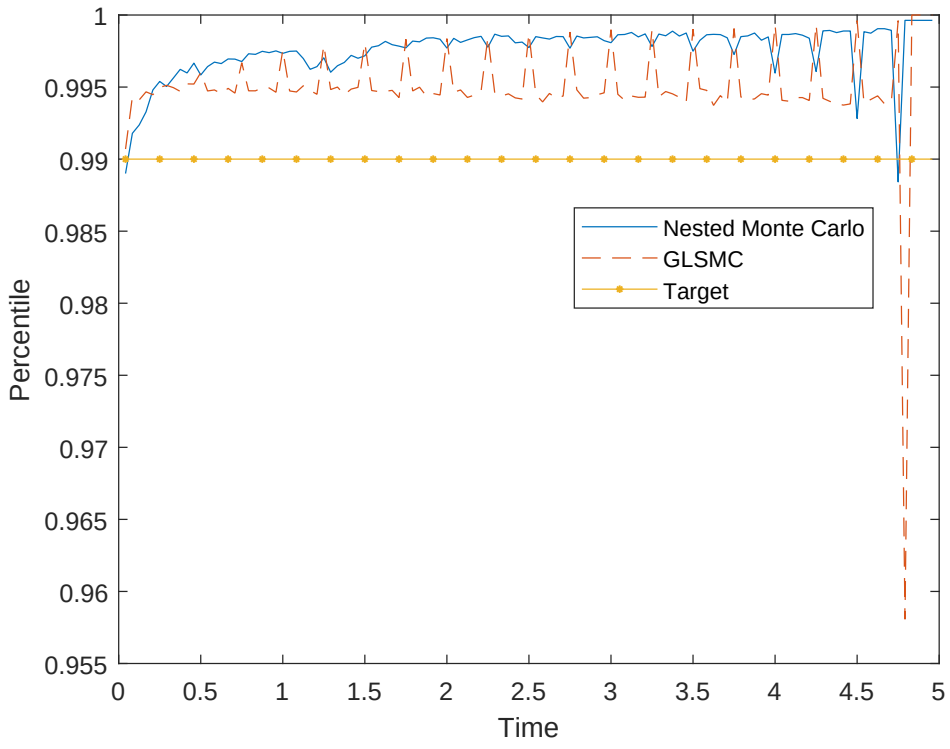

Fig. 5.5: Percentile of simulations with insufficient initial margin.

again, the blue line shows the nested Monte Carlo values, while the orange dashed line indicates the GLSMC values.

As is the case in the other IRS figures, there are spikes at each payment date. In general, initial margin decreases over time, as there are fewer payments to come. The initial margin values between the nested Monte Carlo and GLSMC begin fairly close together, but begin to diverge as the contract nears its completion. In general, the GLSMC spikes up higher leading into payments, and drops down lower following payments — a more extreme movement. It is possible that the nested Monte Carlo values would vary more were the sample sizes used increased, as this could lead to more extreme values being simulated.

To address the results in Figure 5.6, Table 5.1 shows the MVA estimates for the IRS using the two techniques. Similarly to the previous chapter, in the formula

$$\text{MVA}_t = B_t \int_t^T \mathbb{E}^{\mathbb{P}}\left[f_s \frac{\text{IM}_s}{B_s} | \mathcal{F}_t \right] ds, \tag{5.6}$$

the funding spread, f_s, was taken to be 2%. As explained in Section 5.2, it was necessary to simulate jointly the short rate and the integral of the short rate to calculate

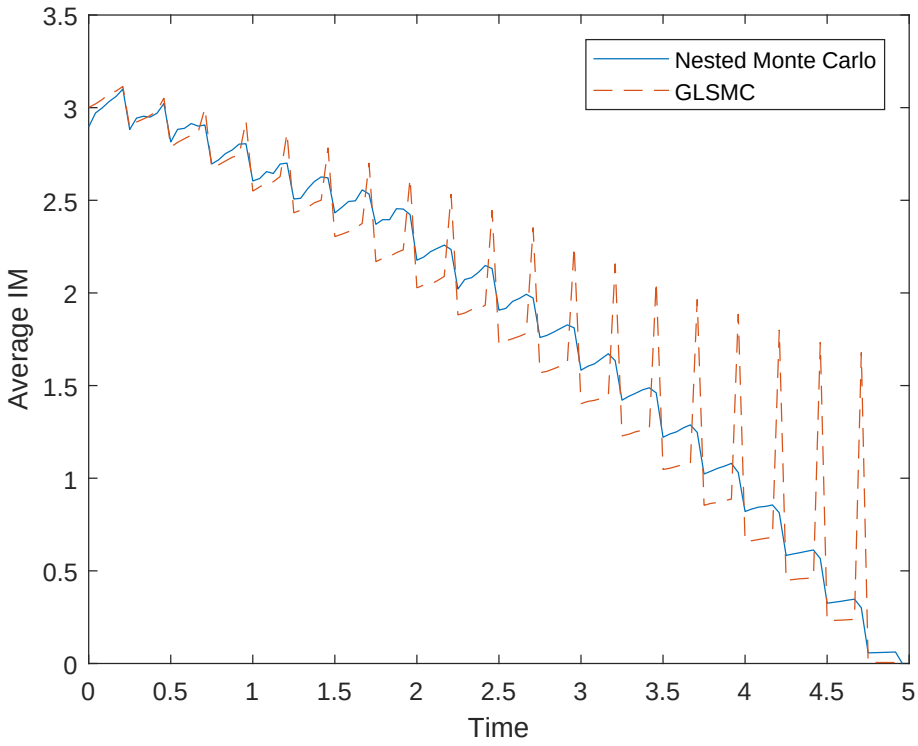

Fig. 5.6: Initial margin over time.

Tab. 5.1: IRS floating-payer MVA estimates.

Technique	MVA
Nested Monte Carlo	0.16183
GLSMC	0.15914

the MVA. The expectation was estimated using the sample means, while the integral was estimated using a trapezoidal technique. The nested Monte Carlo MVA is 1.69% greater than the GLSMC estimate — a minimal difference. This suggests that the points at which the GLSMC IM values spike up in Figure 5.6 are offset cost-wise by the periods where the nested Monte Carlo IM values exceed the GLSMC values.

Figure 5.7 shows the integrands from (5.6) over time. The integrands have a similar shape to the average IM values, but scaled down by the numeraire and the funding spread. Over time, the present value of the expected IM cost is decreasing, with spikes at payment dates. This is expected behaviour, and consistent with previous results.

The MVA values are the additional cost that would need to be accounted for by

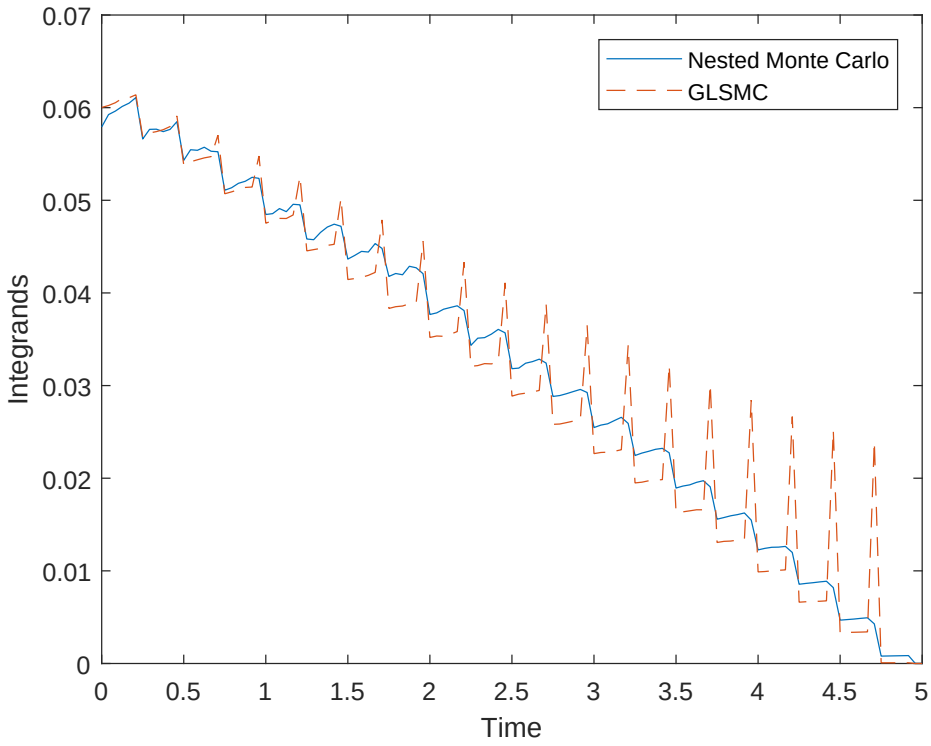

Fig. 5.7: Integrands of MVA.

the party paying the floating payments and receiving the fixed payments. Table 5.2 shows the MVA estimates for the party receiving floating payments and paying the fixed payments. In this case, the nested Monte Carlo estimate is 9.98% smaller than the GLSMC estimate. The GLSMC estimates in both tables are roughly the same, while there exists a significant difference between the nested Monte Carlo values. Since the contract has an initial value of zero, it is to be expected that the MVA costs associated with it are the same for both parties. The difference in the two nested Monte Carlo estimates is likely a result of the limited sample size. Appendix A recreates a number of the figures presented in this chapter from the viewpoint of the party receiving fixed and paying floating, which illustrates the consistency between both sides, and suggests that the GLSMC MVA estimate is the more accurate one. In addition, the GLSMC had a significantly faster runtime, in the region of 40 seconds. The nested Monte Carlo technique took around 450 seconds to run — many times slower than the GLSMC. Addressing the failings of the nested Monte Carlo require increasing sample size to reduce variance, but this comes at the cost of even greater run times.

It is important to note that these calculations were all performed using the aver-

Tab. 5.2: IRS fixed-payer MVA estimates.

Technique	MVA
Nested Monte Carlo	0.14328
GLSMC	0.15916

age initial margin values. While the nested Monte Carlo and GLSMC gave similar results in this case, there arise inconsistencies as one considers extreme portfolio values. For example, nested Monte Carlo often suggests a much higher initial margin for large portfolio values than GLSMC. There are two possible explanations for this: the limited nested Monte Carlo sample size means that the linear interpolation estimation loses accuracy when considering large portfolio values; or, the normality assumption of GLSMC is violated by outliers. The limited sample size of the nested Monte Carlo can also result in negative values for initial margin. This is again due to the linear interpolation failing when working with extreme values. However, these issues do not significantly impact the performance of either technique in addressing PFE, and both provide similar MVA estimates, which supports their accuracy.

Chapter 6

Conclusion

In Chapter 1, we introduced the concept of counterparty exposure and the necessity of its measurement and control, as a result of the financial crisis. Associated with exposure is the concept of potential future exposure, which must be addressed under Basel III. We showed how collateralisation can be implemented in the simple case of using only variation margin, and why this is an insufficient technique as the result of the margin period of risk. An additional form of margining, termed initial margin, was introduced, which led into the next chapter.

Chapter 2 began by detailing the steps involved in calculating exposure as laid out by Pykhtin and Zhu (2007). Following this, the concept of margining was more thoroughly fleshed out, as well as how variation margin, initial margin, and the margin period of risk affect exposure. Finally, numerous valuation adjustments were mentioned, with the margin value adjustment being identified as one of particular interest.

Chapter 3 introduced the two main techniques for calculating initial margin: nested Monte Carlo and Gaussian Least Squares Monte Carlo. The procedures for both were detailed, including the main assumption for the GLSMC (normality of portfolio value changes over the margin period of risk). It was then described how these two techniques would be used in conjunction with independently simulated underlying paths to calculate potential future exposure and thus evaluate the effectiveness of the techniques in estimating initial margin. The previously identified margin value adjustment was given more detail, including how it could be calculated. Lastly, the possible use of the delta and gamma sensitivities to calculate initial margin was explained. The use of this technique was limited to the model studied in the next chapter.

Chapter 4 tested the effectiveness of the various techniques under a geometric Brownian motion model. The option considered was a vanilla FX call option, with numerous assumptions made in order to apply the Black-Scholes formula and sensitivities. It was shown that the nested Monte Carlo provided a good estimate,

albeit at the cost of long computation time. The gamma technique (which used both the gamma and delta of the option) also performed well, but with a tendency to overestimate the initial margin. The GLSMC and delta techniques were partially effective, the delta more so. The normality assumption of the GLSMC was violated by the option and model, especially towards maturity. While not perfect, all of the results showed the effectiveness of using initial margin in reducing PFE.

Chapter 5 considered the more complex G2++ model of Brigo and Mercurio (2006). Under this model, the nested Monte Carlo and GLSMC techniques were tested on a fixed-for-floating interest rate swap. In this scenario, both techniques performed well in reducing PFE. A bit of analysis showed that the normality assumption of the GLSMC technique was better met under the G2++ model than the GBM model. Having confirmed this, subsequent investigation compared various aspects of the initial margin estimates of the two techniques. They were shown to have similar characteristics, and gave comparable estimates for the margin value adjustment.

Bibliography

Andersen, L. B., Pykhtin, M. and Sokol, A. (2016). Credit exposure in the presence of initial margin. Available at SSRN: https://ssrn.com/abstract=2806156 or https://dx.doi.org/10.2139/ssrn.2806156.

Andersen, L. B., Pykhtin, M. and Sokol, A. (2017a). Rethinking the margin period of risk, *Journal of Credit Risk* **13**(1): 1–45.

Andersen, L., Duffie, D. and Song, Y. (2019). Funding value adjustments, *The Journal of Finance* **74**(1): 145–192.

Andersen, L., Pykhtin, M. and Sokol, A. (2017b). Does initial margin eliminate counterparty risk?, *Risk Magazine*, May 2017, pp. 74–79.

Anfuso, F., Aziz, D., Giltinan, P. and Loukopoulos, K. (2017). A sound modelling and backtesting framework for forecasting initial margin requirements, *Risk.net*, May 2017, pp. 86–91.

Basel Committee and others (2014). The standardised approach for measuring counterparty credit risk exposures, *Basel Committee on Banking Supervision*.

Brigo, D. and Mercurio, F. (2006). *Interest Rate Models-Theory and Practice: With Smile, Inflation and Credit*, 2 edn, Springer Finance, Springer: New York.

Canabarro, E. and Duffie, D. (2003). Measuring and marking counterparty risk, *Asset/Liability Management for Financial Institutions*, Euromoney books, pp. 122–134.

International Swap and Derivatives Association (2018). ISDA SIMM methodology. Version 2.1, December.

Lomibao, D. and Zhu, S. (2005). A conditional valuation approach for path-dependent instruments, *in* M. Pyhktin (ed.), *Counterparty Credit Risk Modelling: Risk Management Pricing and Regulation*, Risk Books, chapter 4.

Lopes, M. (2018). *Bias-free joint simulation of multi-factor short rate models and discount factor*, Master's thesis, University of Cape Town.

McWalter, T., Kienitz, J., Nowaczyk, N., Rudd, R. and Acar, S. K. (2018). Dynamic initial margin estimation based on quantiles of Johnson distributions. Available at SSRN: https://ssrn.com/abstract=3147811 or https://dx.doi.org/10.2139/ssrn.3147811.

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

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

Appendix A

Appendix

A.1 Figures for Floating Payer

This appendix recreates a number of the figures from Chapter 5. Here, figures are shown from the perspective of the party receiving fixed and paying floating.

Figure A.1 shows the first 20 simulated IRS value paths over time. This replicates Figure 5.2.

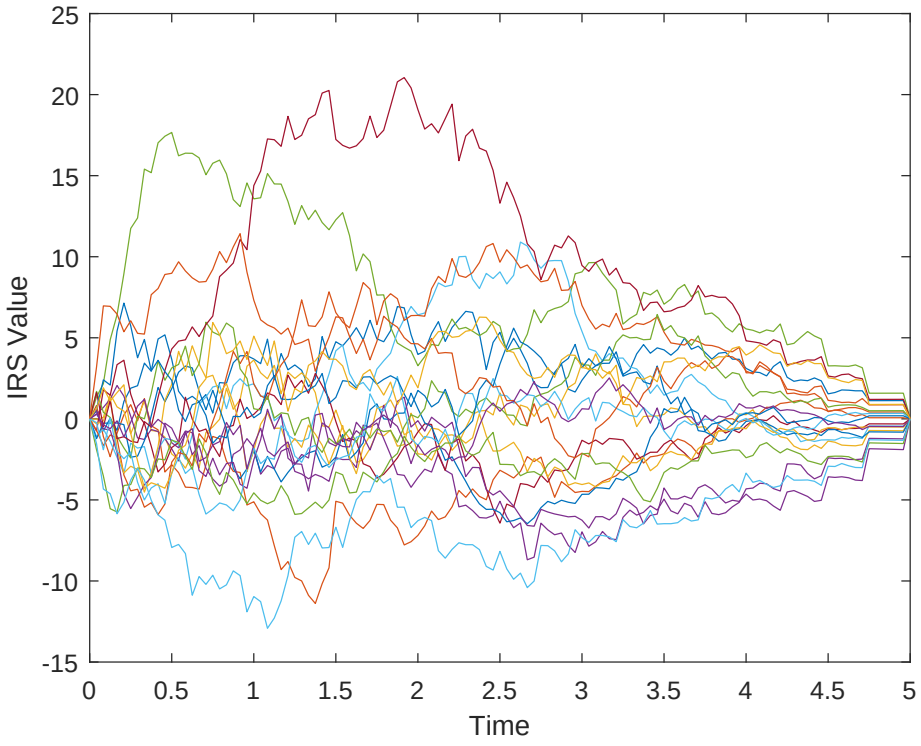

Fig. A.1: IRS value over time.

Figure A.2 shows the PFE of the IRS over time using various collateralisation techniques. This replicates Figure 5.3.

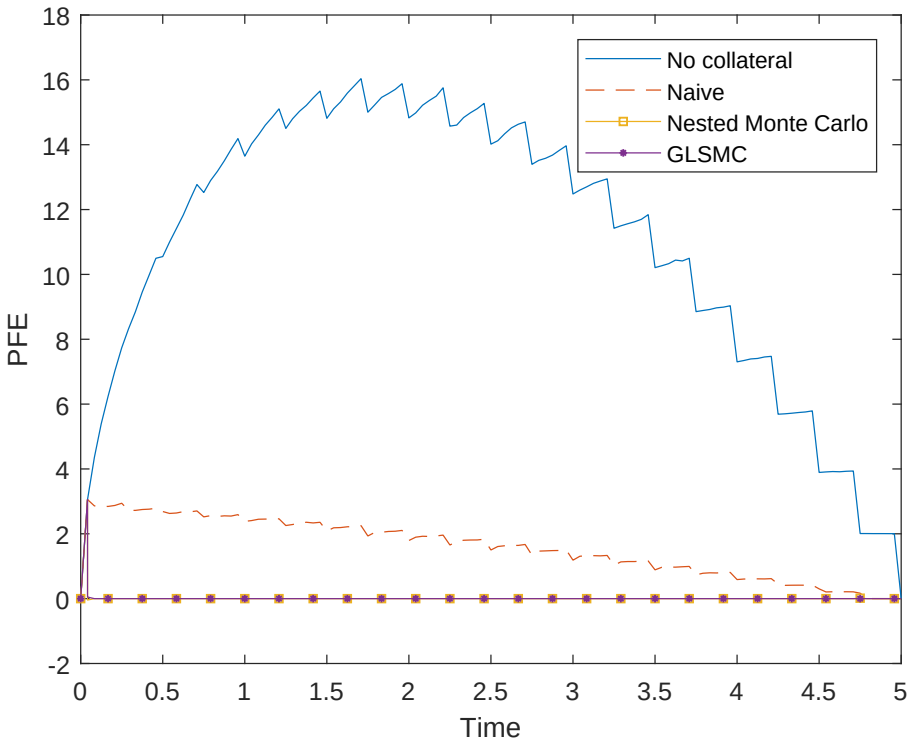

Fig. A.2: IRS PFE over time.

Figure A.3 shows the changes in portfolio value over the margin period of risk halfway through the lifespan of the contract. This replicates Figure 5.4.

Figure A.4 shows the percentile at which the initial margin is proves to be insufficient, using the two techniques. This replicates Figure 5.5.

In all cases, the results for the floating payer are very similar to those of the fixed payer. This is expected, since the IRS is set to have zero value at inception.

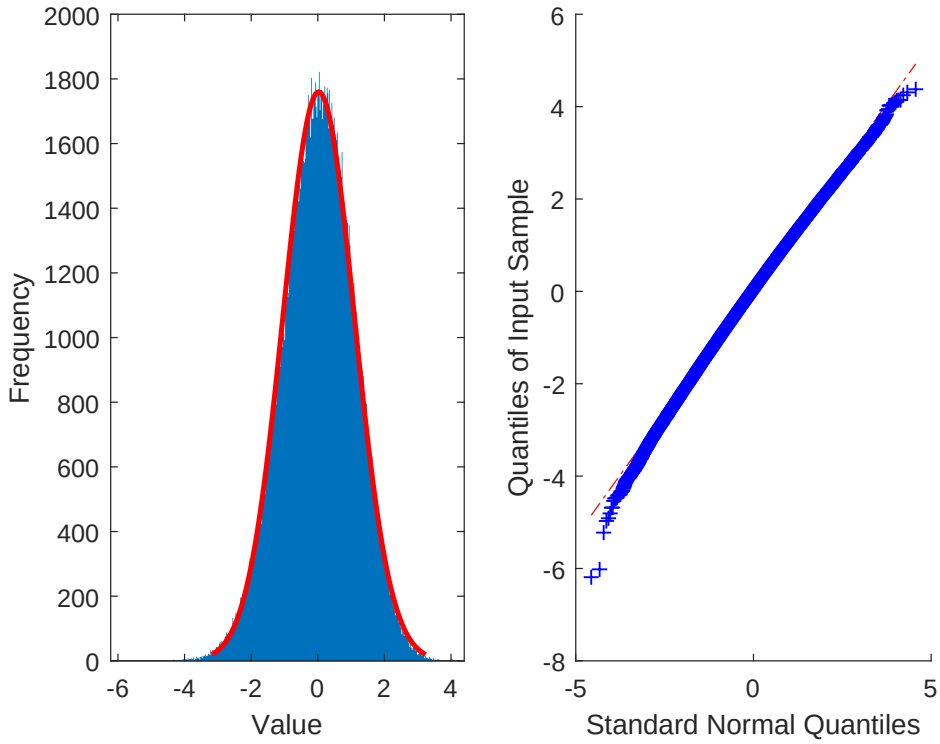

Fig. A.3: IRS value changes over MPOR.

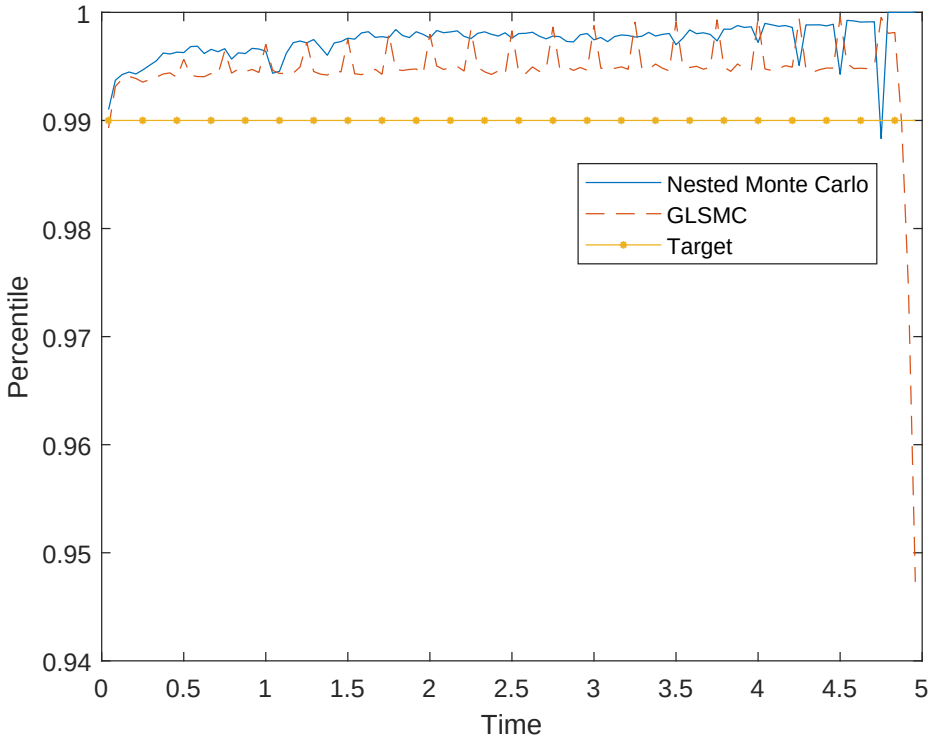

Fig. A.4: Percentile of simulations with insufficient initial margin.