Exposure Modelling Under Change of Measure

Christopher Roberts

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.

March 27, 2017

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 in the University of Cape Town. It has not been submitted before for any degree or examination in any other University.

Signed by candidate

Signature removed

March 27, 2017

Abstract

The credit risk of a portfolio is often managed via measures of counterparty exposure, such as potential future exposure (PFE) and expected exposure (EE), with these measures playing an important role in setting economic and regulatory capital levels. For the sake of risk measurement and risk management these exposure measures should be computed under the real-world probability measure. However, due to the similarity of these exposure calculations to those used in calculating credit valuation adjustments, some have begun to compute them under the risk-neutral measure instead. This is problematic, as the magnitudes of PFEs and EEs differ under different equivalent martingale measures and their associated numéraires.

Working with the Hull-White (HW) model of the short rate, the effect of a change of measure on the PFE and EE profiles of vanilla interest rate swaps and European swaptions is shown under three common measures: the money-market account measure, the T-forward measure and the Linear Gaussian Markovian (LGM) measure. A modified Least Squares Monte Carlo (LSM) algorithm, which allows for substantial computational savings, is then introduced in order to approximate contract level exposures under each of the aforementioned probability measures. Finally, a change of measure is implemented within the modified LSM algorithm in order to approximate exposure profiles under the real-world measure. The modified LSM algorithm is particularly useful for computing exposure profiles of contracts without closed-form valuation formulae, which would otherwise take significantly longer to compute via a standard Monte Carlo approach.

Acknowledgements

I would like to thank my supervisors, Doctor Jörg Kienitz and Associate Professor Thomas McWalter, for their invaluable input and guidance, as well as for their near inexhaustible patience. I must also express my gratitude to my lecturers for their expert tuition: this dissertation would have been much harder to complete without it. Finally, I must thank the class of 2015 for making this such an entertaining and unforgettable year.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In over-the-counter (OTC) derivative contracts, unlike exchange-traded contracts and those settled through a central clearing party (CCP), each party is exposed to the risk of default of its counterparty. In order to manage this risk, one must be able to quantify it. One such quantification is the credit valuation adjustment (CVA), which assigns a price to the risk of counterparty default. This calculation requires, as an input, a measure of the exposure to the counterparty at each possible time of default. These measures of credit exposure differ from those used in credit risk management as they are computed under an equivalent martingale measure (EMM), rather than under the real-world measure.

Counterparty exposure is defined to be the larger of zero and the value of a portfolio of contracts with a counterparty. One can think of it as the cost, at the time of default, of replacing the portfolio that is lost (Canabarro and Duffie, 2003). Other commonly used exposure measures include potential future exposure (PFE), which corresponds to a certain quantile of the distribution of counterparty exposure at a future date; expected exposure (EE), which is the average of counterparty exposure at a future date; and expected positive exposure (EPE), which is the average of EE over time. These, and other measures, play an important role in risk management and setting of economic and regulatory capital (Stein, 2014).

Calculating the distribution of the exposure to a given counterparty at a future date usually consists of three steps, as described by Pykhtin and Zhu (2007). First, risk factors (such as interest rates, stock prices etc.) which are relevant to the valuation of the contracts are simulated. Having simulated the relevant risk factors the contracts themselves are then valued. In the case of contracts without closed-form valuation formulae this requires an additional set of simulations at each time step in order to value the contracts. Finally, for each realisation of the risk factors, the contract-level exposures are aggregated, subject to any netting or margining agreements, yielding a distribution of counterparty exposure.

For the purposes of risk management and the calculation of economic and regulatory capital, a firm should compute its exposure measures under the real-world measure. Specifically the evolution of the portfolio risk factors should be carried out under the real-world measure, whereas the contract valuation should be performed under an EMM (Canabarro and Duffie, 2003; Pykhtin and Zhu, 2007; Stein, 2014). However, some have taken to calculating exposures completely under a chosen EMM, which implicitly assumes that doing so will not have a material effect on the exposure values. This seems to be motivated by the similarity of CVA calculations to exposure calculations (as CVA calculations require credit exposures under an EMM), as well as by the existence of EMM valuation methodologies and computational architectures (Stein, 2014).

Contract level exposures calculated under an EMM, however, differ not only from those calculated under the real-world measure, but also those calculated under EMMs corresponding to different numéraires. In fact Stein (2014) shows how one can carefully construct a numéraire to give any desired PFE profile when working under its associated EMM.

This dissertation quantifies the impact of different probability measures on exposure calculations. This is achieved by calculating two commonly used exposure measures, namely PFE and EE, for vanilla interest rate swaps and European swaptions under commonly used pairs of EMMs and numéraires. It is shown that the magnitude of the exposure profiles of these basic contracts may differ quite substantially depending on the chosen probability measure.

This dissertation also attempts to resolve the problem of exposures computed under an EMM in two parts. Firstly, a modified Least Squares Monte Carlo (LSM) algorithm is implemented, which avoids the computationally intensive nested Monte Carlo situation that arises when modelling the exposures of contracts without closed-form valuation formulae. Secondly, a change of measure modification is made to the LSM algorithm which allows for the use of real-world risk factor realisations, thus giving rise to real-world, rather than risk-neutral, exposure profiles.

In Chapter 2 a brief overview of the change of numéraire methodology is given. In Chapter 3 the Hull-White short rate model and its generalisation, the Linear Gaussian Markovian model, are described, and the pricing formulae for swaps and swaptions under these models are given. Chapter 4 defines the exposure measures used in this dissertation, outlines the simulation methodology used in calculating swap and swaption exposures, and presents and discusses the resulting exposure profiles under a selection of probability measures. In Chapter 5 the modified LSM algorithm and its change of measure extension are presented, and the accuracy as well as computation times for the approximated exposure profiles are benchmarked against those obtained previously. Chapter 6 concludes the dissertation.

Chapter 2

Change of Numéraire

We provide a brief overview of the change of numéraire technique originally due to Geman *et al.* (1995).

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ be a filtered probability space, with the filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions (i.e. $\mathcal{F}_t$ is right continuous and $\mathcal{F}_0$ contains all the null-sets of $\mathbb{P}$). Let $W_t^{\mathbb{P}}$ be an $\mathcal{F}_t$-adapted d-dimensional Brownian motion under the probability measure $\mathbb{P}$. The measure $\mathbb{P}$ is taken to be the real-world probability measure. This setup forms the base of a multidimensional continuous-time market model in which asset prices, S_t, are modelled by Itō diffusions. That is to say that asset price dynamics are given by

$$dS_t = \mu(t, S_t)\, dt + \sigma(t, S_t)\, dW_t^{\mathbb{P}} \tag{2.1}$$

under $\mathbb{P}$, where $\mu(t, S_t)$ is the drift, and $\sigma(t, S_t)$ the d-dimensional volatility vector of the diffusion process.

Let T be a finite time horizon. A numéraire, N_t, is a non-dividend paying asset which is almost surely strictly positive on $[0, T]$. Often, a numéraire can be thought of as a unit of account, so that if S_t is another asset, then S_t/N_t is the value of S_t in terms of units of N_t. A common numéraire is the money-market account, A_t, which corresponds to an investment of one unit of currency at time 0 in an account that grows at the continuously compounded rate of interest r_t. We can then interpret S_t/A_t as the discounted value of S_t.

If X_T is an attainable claim (i.e. there exists a self-financing portfolio of assets which replicates X_T in all states of the world), and N_t a valid numeráire, then

$$\frac{X_t}{N_t} = \mathrm{E}_{\mathbb{Q}^N}\left[\left.\frac{X_T}{N_T}\right| \mathcal{F}_t\right], \tag{2.2}$$

where $\mathbb{Q}^N$ is an equivalent martingale measure (EMM) for N_t. We call a measure $\mathbb{Q}^N$ an EMM if $\mathbb{Q}^N(F) = 0 \iff \mathbb{P}(F) = 0\ \forall F \in \mathcal{F}$ (i.e. $\mathbb{Q}^N$ and $\mathbb{P}$ agree on the null sets of $\mathcal{F}$) and if the asset ratios S_t/N_t are martingales under $\mathbb{Q}^N$. A fundamental EMM is the measure associated with the money-market account, A_t, which is often denoted simply by $\mathbb{Q}$. This measure is usually referred to as *the* risk-neutral measure.

If $\hat{N}_t$ is another valid numéraire, then we have

$$N_t \mathrm{E}_{\mathbb{Q}^N}\left[\left.\frac{X_T}{N_T}\right|\mathcal{F}_t\right] = \hat{N}_t \mathrm{E}_{\mathbb{Q}^{\hat{N}}}\left[\left.\frac{X_T}{\hat{N}_T}\right|\mathcal{F}_t\right], \tag{2.3}$$

where $\mathbb{Q}^{\hat{N}}$ is an EMM for $\hat{N}_t$. It can be shown that the Radon-Nikodym derivative that effects the change of measure from $\mathbb{Q}^N$ to $\mathbb{Q}^{\hat{N}}$ is given by

$$\frac{d\mathbb{Q}^{\hat{N}}}{d\mathbb{Q}^N} = \frac{\hat{N}_T/N_T}{\hat{N}_0/N_0}. \tag{2.4}$$

As a consequence of this we can obtain asset price dynamics under $\mathbb{Q}^{\hat{N}}$ from dynamics under $\mathbb{Q}^N$ by means of a Girsanov transformation, the kernel of which is the volatility, $\hat{\sigma}(t, \hat{N}_t)$, of the new numeráire $\hat{N}_t$. Thus, an Itō diffusion with dynamics

$$dS_t = \mu(t, S_t)\,dt + \sigma(t, S_t)\,dW_t^{\mathbb{Q}^N} \tag{2.5}$$

under $\mathbb{Q}^N$, where $W_t^{\mathbb{Q}^N}$ is a $\mathbb{Q}^N$ Brownian motion, will have the following dynamics under $\mathbb{Q}^{\hat{N}}$

$$dS_t = \left(\mu(t, S_t) + \sigma(t, S_t)\hat{\sigma}(t, \hat{N}_t)^\top\right)\,dt + \sigma(t, S_t)\,dW_t^{\mathbb{Q}^{\hat{N}}}, \tag{2.6}$$

where $W_t^{\mathbb{Q}^{\hat{N}}}$ is a $\mathbb{Q}^{\hat{N}}$ Brownian motion.

Chapter 3

Interest Rate Models

For the purposes of this dissertation two interest rate models are considered, namely the short rate model due to Hull and White (1990) and the Linear Gaussian Markovian model of Hagan and Woodward (1999).

3.1 Hull-White Model

In the Hull-White Extended Vasiček model, which this dissertation refers to as simply the Hull-White (HW) model, the short rate has the following dynamics under the risk-neutral measure $\mathbb{Q}$:

$$dr_t = (\theta(t) - ar_t)\,dt + \sigma\,dW_t^{\mathbb{Q}}, \tag{3.1}$$

where θ is a deterministic function of time, and $W_t^{\mathbb{Q}}$ is a $\mathbb{Q}$ Brownian motion. The function $\theta(t)$ allows the model to fit the initial term structure of interest rates exactly with

$$\theta(t) = f_T^M(0,t) + f^M(0,t) + \frac{\sigma^2}{2a}\left(1 - e^{-2at}\right). \tag{3.2}$$

Here $f^M(0,t)$ is the market-observed instantaneous forward rate for maturity t, and $f_T^M(0,t)$ the partial derivative of $f^M(0,t)$ with respect to its second argument. Integrating (3.1) we find that

$$r_t = r_s e^{-a(t-s)} + \alpha(t) - \alpha(s)e^{-a(t-s)} + \sigma\int_s^t e^{-a(t-u)}\,dW_u^{\mathbb{Q}}, \tag{3.3}$$

where

$$\alpha(t) = f^M(0,t) + \frac{\sigma^2}{2a^2}\left(1 - e^{-at}\right)^2. \tag{3.4}$$

Thus, conditional on $\mathcal{F}_s$, r_t is normally distributed under $\mathbb{Q}$ with mean and variance:

$$\mathbb{E}[r_t \mid \mathcal{F}_s] = r_s e^{-a(t-s)} + \alpha(t) - \alpha(s)e^{-a(t-s)} \tag{3.5}$$

$$\mathrm{Var}[r_t \mid \mathcal{F}_s] = \frac{\sigma^2}{2a}\left(1 - e^{-2a(t-s)}\right). \tag{3.6}$$

Since the HW model is an affine term structure model, zero-coupon bond prices have the following form:

$$\begin{aligned} P(t,T) &= A(t,T)e^{-B(t,T)r_t} \\ A(t,T) &= \frac{P^M(0,T)}{P^M(0,t)} \exp\left\{B(t,T)f^M(0,t) - \frac{\sigma^2}{4a}\left(1-e^{-2at}\right)B(t,T)^2\right\} \\ B(t,T) &= \frac{1}{a}\left(1-e^{-a(T-t)}\right), \end{aligned} \tag{3.7}$$

with $P(t,T)$ being the value of a zero-coupon bond at time t which matures at T, and $P^M(t,T)$ the market-observable zero-coupon bond price.

Under the T-forward measure we take the numéraire to be the zero-coupon bond maturing in T years time. An application of Itō's lemma to the bond price given by (3.7) shows that this bond has volatility $-\sigma B(t,T)$. Thus the dynamics of the short rate under the T-forward measure are

$$dr_t = \left(\theta(t) - ar_t - \sigma^2 B(t,T)\right)\,dt + \sigma\,dW_t^{\mathbb{Q}^T}, \tag{3.8}$$

where $W_t^{\mathbb{Q}^T}$ is a $\mathbb{Q}^T$ Brownian motion. Under the T-forward measure the conditional mean of r_t is

$$\begin{aligned} \mathbb{E}[r_t \mid \mathcal{F}_s] = r_s e^{-a(t-s)} - \frac{\sigma^2}{a^2}\left(1-e^{-a(t-s)}\right) + \frac{\sigma^2}{2a^2}\left(e^{-a(T-t)} - e^{-a(T+t-2s)}\right) \\ + \alpha(t) - \alpha(s)e^{-a(t-s)}, \end{aligned} \tag{3.9}$$

where $\alpha(t)$ is defined as before. The conditional variance remains the same as that given in (3.6).

3.2 Linear Gaussian Markovian Model

Rather than attempt to model some idealised quantity such as the short rate or instantaneous forward rate, Hagan and Woodward (1999) instead choose to specify their model in terms of an arbitrary state process X_t, a general numéraire $N(t, X_t)$, and the martingale valuation formula. We follow the exposition of the Linear Gaussian Markovian (LGM) model given in Hagan (2009). The state process has dynamics

$$dX_t = \alpha(t)\,dW_t\,, \tag{3.10}$$

where $X_0 = 0$, and the numéraire is given by

$$\begin{aligned} N(t, X_t) &= \frac{1}{P^M(0,t)} e^{H(t)X_t + \frac{1}{2}H(t)^2\zeta(t)} \\ \zeta(t) &= \int_0^t \alpha(s)^2\,ds\,. \end{aligned} \tag{3.11}$$

Here $\alpha(t)$ and $H(t)$ are deterministic, time-varying parameters. It can be shown that the LGM model is equivalent to the HW model when we set

$$H(t) = e^{-at},\ \alpha(t) = \frac{\sigma}{a}e^{at}. \tag{3.12}$$

Due to the specification of the numéraire, the LGM model automatically recovers the initial zero-coupon bond curve, and it can be shown that zero-coupon bond prices are given by

$$P(t,T) = \frac{P^M(0,T)}{P^M(0,t)}e^{(H(t)-H(T))X_t+\frac{1}{2}(H^2(t)-H^2(T))\zeta(t)}. \tag{3.13}$$

3.3 Pricing Interest Rate Instruments

3.3.1 Vanilla Interest Rate Swaps

A vanilla interest rate swap (IRS) is a contract between two counterparties where one counterparty agrees to pay a series of fixed interest rate payments (the fixed leg) and receive a series of floating interest rate payments (the floating leg) from the second counterparty. The floating cashflows are based on a floating reference rate (such as 3 month LIBOR) which is fixed one period before each payment date in order to calculate the payment due on the floating leg. From the point of view of one counterparty, a payer swap is one in which the fixed leg is paid in exchange for the floating leg, whereas in a receiver swap the floating leg is paid in exchange for the fixed leg.

An IRS, based on a notional N, which is traded at T_0 and which matures at T_N, will have a series of reset dates $T_0, T_1, ..., T_{N-1}$ on which the floating rate is fixed, and a series of payment dates $T_1, T_2, ..., T_N$ on which cashflows are exchanged. It can be shown that the value of a payer swap contract at any $T_0 \leq t < T_N$ is:

$$V_{\text{IRS}}(t) = N\left((1 + L(T_{k-1}, T_k)\tau_k)P(t, T_k) - P(t, T_N) - \sum_{i=k}^{N} P(t, T_i)K\tau_i\right), \tag{3.14}$$

where $k = \max\{i : T_{i-1} \leq t\}$ is the index of the most recent reset date before t. Here $L(T_{k-1}, T_k)$ is the most recent fixing of the simple floating rate at T_{k-1}, K is the fixed rate and $\tau_i = T_i - T_{i-1}$ is the accrual period between the reset date T_{i-1} and payment date T_i.

The fair swap rate (FSR) is the fixed rate which sets the value of an IRS to zero at T_0. One can show that this is given by

$$\text{FSR} = \frac{1 - P(T_0, T_N)}{\sum_{i=1}^{N} P(T_0, T_i)\tau_i}. \tag{3.15}$$

3.3.2 European Swaptions in The HW Model

An ordinary European swaption is an option which offers the holder the right to enter into a swap at maturity of the option, at a predetermined swap rate (the strike rate). A payer swaption (PS) is an option on a payer IRS, and a receiver swaption (RS) is an option on a receiver IRS.

The valuation of European swaptions under the HW model is outlined in Brigo and Mercurio (2007). Consider a payer swaption, with notional N, which matures at some future time T_0 and has a strike rate K. Further suppose that the first reset of the underlying swap coincides with T_0, and that the underlying swap matures at T_n. This swaption has the following payoff at T_0:

$$N\left(1 - P(T_0, T_N) - \sum_{i=1}^{n} K\tau_i P(T_0, T_i)\right)^{+} . \tag{3.16}$$

Through an application of Jamshidian's trick (Jamshidian, 1989), this swaption can be valued at $t < T_0$ as an option on a coupon bearing bond. Let $c_i = K\tau_i$ for $i = 1, 2, ..., n-1$ and $c_i = 1 + K\tau_i$ for $i = n$. First, an r^* must be found such that

$$\sum_{i=1}^{n} c_i A(T_0, T_i) e^{-B(T_0,T_i)r^*} = 1\,. \tag{3.17}$$

Then, setting $K_i = A(T_0, T_i)e^{-B(T_0,T_i)r^*}$, the value of the payer swaption at t is given by

$$V_{\text{PS}}(t) = N \sum_{i=1}^{n} c_i \text{ZBP}(t; T_0, T_i, K_i)\,, \tag{3.18}$$

where $\text{ZBP}(t; T_0, T_i, K_i)$ is the time t value of a put maturing at T_0 on zero-coupon bond with maturity T_i, struck at K_i. Each zero-coupon bond put has value

$$\text{ZBP}(t, T_0, T_i, K_i) = K_i P(t, T_0)\Phi(-h_i + \sigma_{p,i}) - P(t, T_i)\Phi(-h_i)\,, \tag{3.19}$$

where Φ is the standard normal cumulative distribution function, and

$$\begin{aligned} \sigma_{p,i} &= \sigma\sqrt{\frac{1 - e^{-2a(T_0 - t)}}{2a}} B(T_0, T_i) \\ h_i &= \frac{1}{\sigma_{p,i}} \ln\left(\frac{P(t, T_i)}{P(t, T_0)K_i}\right) + \frac{\sigma_{p,i}}{2}. \end{aligned} \tag{3.20}$$

Similarly the value of a receiver swaption is

$$V_{\text{RS}}(t) = N \sum_{i=1}^{n} c_i \text{ZBC}(t; T_0, T_i, K_i)\,, \tag{3.21}$$

where $\text{ZBC}(t; T_0, T_i, K_i)$ is a zero-coupon bond call with value

$$\text{ZBC}(t, T_0, T_i, K_i) = P(t, T_i)\Phi(h_i) - K_i P(t, T_0)\Phi(h_i - \sigma_{p,i})\,, \tag{3.22}$$

and $\sigma_{p,i}$ and h_i are defined as before.

3.3.3 European Swaptions in The LGM Model

Due to the specification of the LGM model in terms of a Gaussian state process, X_t, and numéraire, $N(t, X_t)$, claims can be valued directly by integration. That is, given a contract with payoff $V(T, X_T)$ at maturity its value today is given by

$$V(t, X_t) = N(t, X_t) \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi(\zeta(T) - \zeta(t))}} e^{-\frac{(x-X_t)^2}{2(\zeta(T)-\zeta(t))}} \frac{V(T, x)}{N(T, x)} \, dx \,. \tag{3.23}$$

In the case of a payer swaption, struck at K and maturing at T_0, the following numéraire-deflated payoff must be integrated:

$$\Bigg[P(0, T_0)e^{-H(T_0)x-\frac{1}{2}H(T_0)^2\zeta(T_0)} - \sum_{i=1}^{n} \tau_i K P(0, T_i)e^{-H(T_i)x-\frac{1}{2}H(T_i)^2\zeta(T_0)} - P(0, T_n)e^{-H(T_n)x-\frac{1}{2}H(T_n)^2\zeta(T_0)}\Bigg]^+ . \tag{3.24}$$

To evaluate (3.23) in this instance one must integrate over the region where the payoff is non-zero. Because $H(t) = e^{-at}$ in the Hull-White equivalent formulation of the LGM model, $H(T_i) - H(T_0)$ is negative for each i and thus one can find a unique x^* such that

$$P(0, T_0) = \sum_{i=1}^{n} \tau_i K P(0, T_i)e^{-(H(T_i)-H(T_0))x^*-\frac{1}{2}(H(T_i)^2-H(T_0)^2)\zeta(T_0)} + P(0, T_n)e^{-(H(T_n)-H(T_0))x^*-\frac{1}{2}(H(T_n)^2-H(T_0)^2)\zeta(T_0)}, \tag{3.25}$$

and for which the integrand is positive for $x < x^*$. Substituting the numéraire deflated payoff (3.24) into (3.23) and integrating yields

$$\begin{aligned} V_{\text{PS}}(t, X_t) = {} & \frac{P(0, T_0)}{P(0, t)} e^{-(H(T_0)-H(t))X_t-\frac{1}{2}(H(T_0)^2-H(T_t)^2)\zeta(t)} \Phi(d_0) - \\ & \sum_{i=1}^{n} \tau_i K \frac{P(0, T_i)}{P(0, t)} e^{-(H(T_i)-H(t))X_t-\frac{1}{2}(H(T_i)^2-H(t)^2)\zeta(t)} \Phi(d_i) \\ & + \frac{P(0, T_n)}{P(0, t)} e^{-(H(T_n)-H(t))X_t-\frac{1}{2}(H(T_n)^2-H(t)^2)\zeta(t)} \Phi(d_n) \,, \end{aligned} \tag{3.26}$$

where Φ is the standard normal cumulative density function and d_i is given by

$$d_i = \frac{(x^* - X_t) + H(T_i)(\zeta(T_0) - \zeta(t))}{\sqrt{\zeta(T_0) - \zeta(t)}} \,. \tag{3.27}$$

The value of a receiver swaption can be determined by put-call parity to be

$$V_{\text{RS}}(t, X_t) = V_{\text{PS}}(t, X_t) + N(t, X_t) \left(\sum_{i=1}^{n} \tau_i K P(0, T_i) + P(0, T_n) - P(0, T_0) \right) . \tag{3.28}$$

Chapter 4

Exposure Calculations

4.1 Exposure Measures

Canabarro and Duffie (2003) give basic definitions for a variety of exposure measures, which are elaborated upon by Stein (2014). Let V_t be the value of a portfolio, without netting or collateral agreements, at time t. Then the p% PFE at time horizon T is given by X where

$$\mathbb{P}[V_T < X] = \mathbb{E}^{\mathbb{P}}[I_{V_T<X}] = p\,. \tag{4.1}$$

This is analogous to VaR as it gives a quantile which is not exceeded $(1-p)$% of the time. Noting that the actual exposure at time t is given by $\max(V_t, 0)$, we have that the EE at time T is

$$\text{EE}(V,T) = \mathbb{E}^{\mathbb{P}}[\max(V_T, 0)] = \mathbb{E}^{\mathbb{P}}[V_T I_{V_T \geq 0}]\,. \tag{4.2}$$

Similarly, the EPE between times t_1 and t_2 is

$$\text{EPE}(V, t_1, t_2) = \frac{1}{t_2 - t_1}\int_{t_1}^{t_2} \text{EE}(V, s)\, ds\,. \tag{4.3}$$

The effective expected exposure (EEE) and effective expected positive exposure (EEPE) are defined similarly in terms of the maximum EE over the period t_1 to t_2:

$$\text{EEE}(V, t_1, t_2) = \max_{t_1 \leq s \leq t_2} (\text{EE}(s)) \tag{4.4}$$

$$\text{EEPE}(V, t_1, t_2) = \frac{1}{t_2 - t_1}\int_{t_1}^{t_2} \text{EEE}(V, t_1, s)\, ds\,. \tag{4.5}$$

For the purposes of this dissertation only the PFE and EE exposure measures are considered. The other measures given above do, however, play a role in the measuring of counterparty credit exposure, and in determining capital requirements (Stein, 2014).

4.2 Simulation Methodology

We consider two types of interest rate derivative contracts in order to determine what effect a change of measure has on exposure profiles. The first is a 10 year vanilla IRS, struck at the fair swap rate. The second is a 1 year into 5 year swaption, struck at 0.05%. For each type of contract we consider both the payer and receiver versions.

Both the HW and LGM models are parameterised in terms of an initial term structure of forward rates, or equivalently zero-coupon bond prices. This is fixed to be:

$$\begin{aligned} f(0,T) &= 0.1 - 0.05e^{-0.0458T} \\ P(0,T) &= \exp\left(-0.1T - \frac{250}{229}(e^{-0.0458T} - 1)\right) . \end{aligned} \tag{4.6}$$

For the HW model we choose $a = 0.05$ and $\sigma = 0.01$, with these being considered "reasonable" parameter values. The LGM parameters $H(t)$ and $\alpha(t)$ are then determined by the equations in (3.12).

To obtain the exposure profiles of each contract we proceed as follows:

1. Discretise the time grid $[0, T]$, where T is the contract's maturity, into n fixed time steps of length $\Delta T = T/n$.

2. Simulate N independent paths of the underlying risk factor at time horizons $\Delta T, 2\Delta T, ..., n\Delta T = T$.

3. Calculate the contract's value, along each path, at each time horizon.

4. Compute the desired exposure measure at each time horizon to arrive at the contract exposure profile.

For our purposes we choose $N = 5000$ for both the swap and swaption exposure calculations. For the swap calculations we choose a time step of $\Delta T = 0.0625$, and for the swaptions we take $\Delta T = 0.025$. We calculate the exposures under four different measures: the money-market account (MMA) measure, 20-Forward Measure, 60-Forward Measure and the LGM measure. For both the swaps and swaptions we work with a unit notional.

It is worth noting that since the risk factors r_t and X_t are Gaussian, they could be simulated directly, instead of having to apply an Euler scheme to their respective SDEs. Moreover, since the risk factors are Gaussian, and since each contract's value is monotonic in its underlying risk factor, one could instead compute an exact p% PFE by valuing the contract at the corresponding percentile of its underlying risk factor. In principal, one might also be able to obtain closed-form expressions for the EE of each contract by direct integration, but this may be considerably more difficult. We choose the Monte Carlo approach, however, as we require risk factor paths and numéraire paths for the LSM approximations, and because we also wish to benchmark these approximations against the Monte Carlo PFEs and EEs.

4.3 Swap Exposure

There are two qualitative features of the swap exposure profiles in Figures (4.1) and (4.2) which are worth noting. The first is that at initiation and maturity the swap exposure is zero, since it was traded at the fair swap rate. Had the swap been traded at a different rate, its PFE and EE at initiation would simply be its net present value. The second is the distinctive sawtooth exposure pattern (more pronounced in the PFEs than the EEs) which is a result of the intermediate swap cashflows. Also worth noting is that the PFE peaks about one-third to one-half of the way through the life of the contract, which is consistent with the literature (Pykhtin and Zhu, 2007).

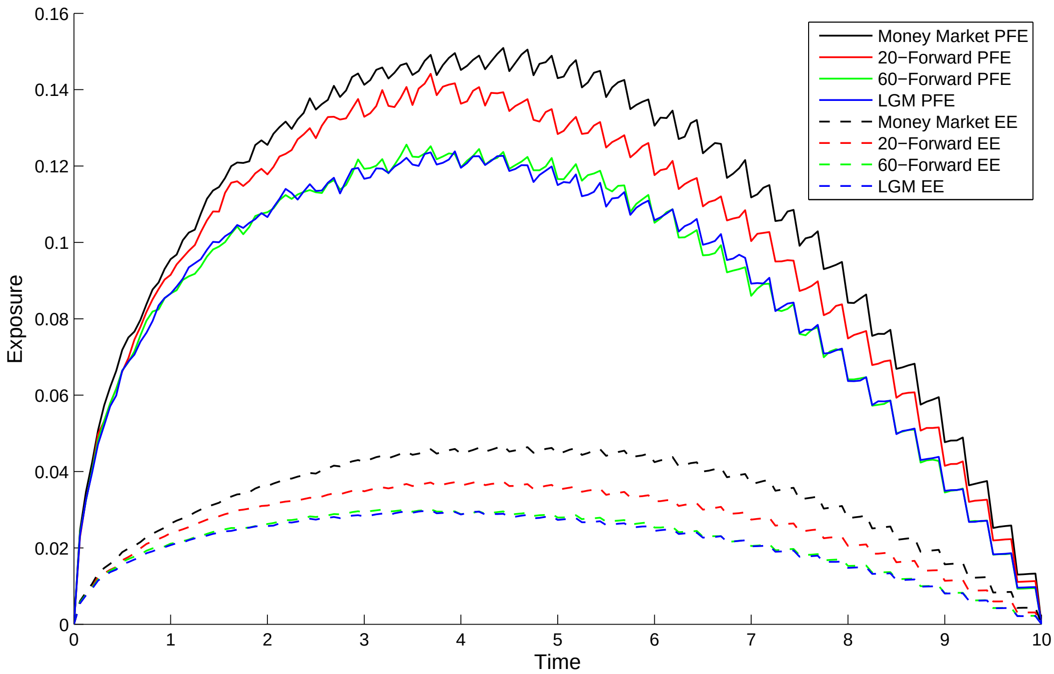

Fig. 4.1: 95% PFE and EE for a 10 year vanilla payer swap.

One can see quite clearly, in both the payer and receiver cases, the effect of calculating contract exposures under different EMMs, with there being as much as 25% difference between the PFE computed under the MMA measure and the PFE

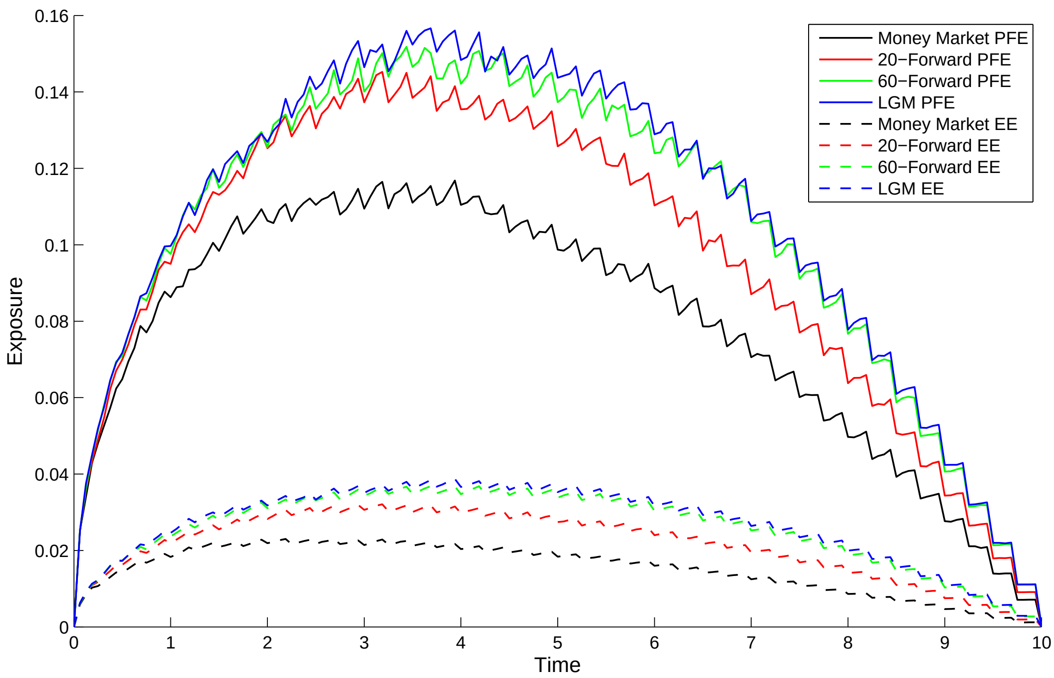

Fig. 4.2: 95% PFE and EE for a 10 year vanilla receiver swap.

computed under the LGM measure for the payer swap (and about a 40% difference in the opposite direction for the receiver swap).

The difference in magnitudes of the exposure profiles under the different measures is simply due to the different drift rates of the risk factors. This can be seen when comparing the MMA, 20-Forward and 60-Forward exposure profiles. Since moving from the MMA measure to a T-forward measure involves a drift adjustment of $-\sigma^2 B(t,T)$, the simulated short rates tend to be lower under the 20-Forward measure than the MMA measure, and lower still under the 60-Forward measure. As a result the PFE of a payer swap is much higher under the MMA measure than under the 20-Forward and 60-Forward measures, since the value of a payer swap is positively related to the short rate. Likewise the PFE of a receiver swap is lower under the MMA measure than under these two forward measures because receiver swap values are negatively related to the short rate. Moreover, since the drift of the short rate remains positive (even after making the adjustment for a change of measure) the simulated short rates tend to drift upwards and thus the payer swap PFEs and EEs peak later than they do for a receiver swap (approximately 4 years for the payer swap, as opposed to about 3 years for the receiver swap).

Although payer (receiver) swap values have a similar positive (negative) relationship with the state process X_t, we cannot appeal directly to this drift argument to explain the difference in exposure values (as X_t is not comparable with r_t under the MMA measure, in the same way that r_t under the MMA measure is comparable with r_t under different measures). We simply remark that the payer and receiver profiles seem similar to those obtained under the 60-Forward measure.

4.4 Swaption Exposure

The effect of a change of measure on contract exposure profiles is even more pronounced in the swaption example than in the swap example. In this instance the LGM PFEs are significantly larger than the PFE profiles under the other measures, with the LGM PFE growing to be as much as 1.65 times greater than the PFEs under the other measures in the payer case (Figure 4.3), and as much as 3.5 times greater in the receiver case (Figure 4.4). The significant differences between the PFEs under the LGM measure and the PFEs under the other measures highlight the problem of computing exposures under an EMM: all four measures are valid, in the sense that they correspond to valid numéraires, yet exposure profiles computed under each measure can vary substantially from those computed under the other measures.

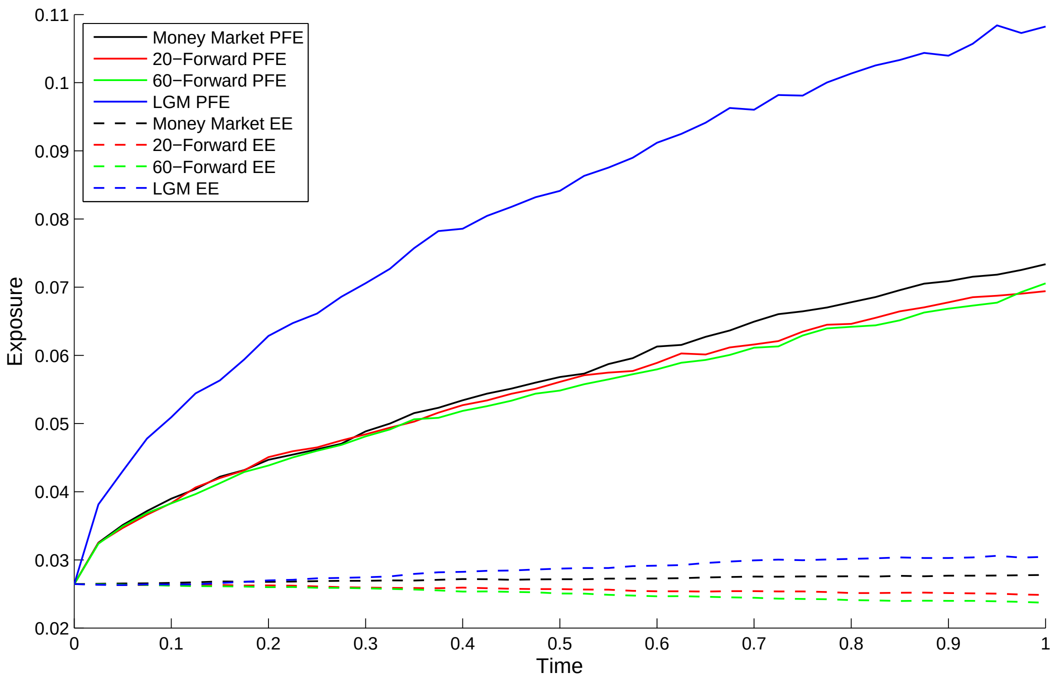

Fig. 4.3: 95% PFE and EE for a 1Y5Y year payer swaption.

The effect of decreasing the drift of the short rate as one moves from the MMA

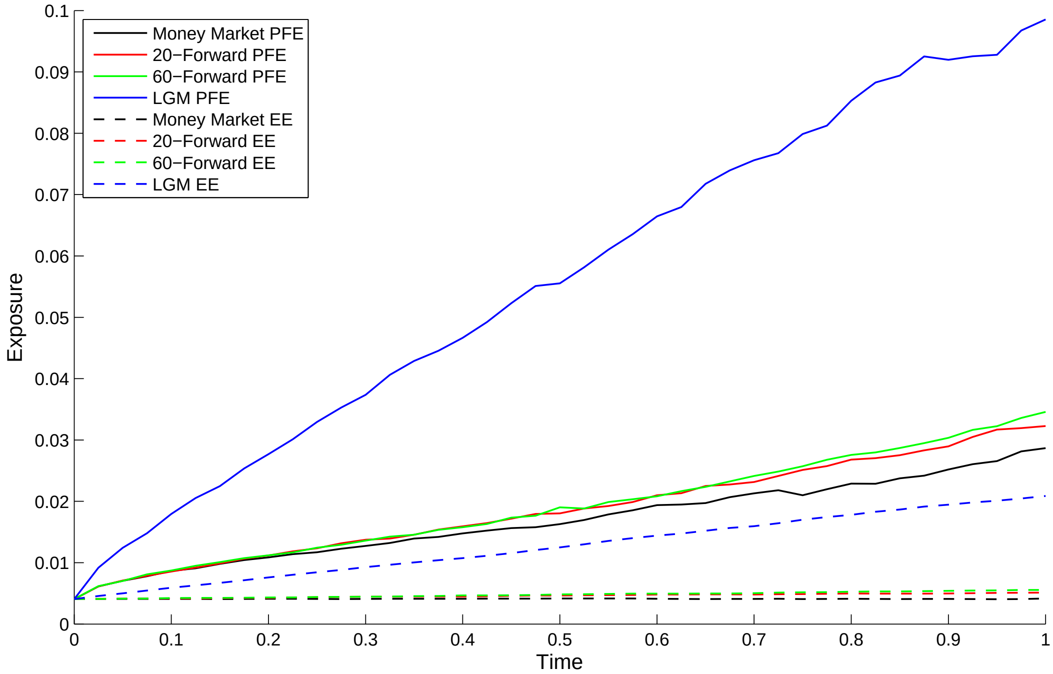

Fig. 4.4: 95% PFE and EE for a 1Y5Y year receiver swaption.

measure to the 20-Forward measure can again be seen in the ordering of their respective PFEs, with the MMA PFE lying above the 20-Forward PFE, which in turn is above the 60-Forward PFE in the payer case and this being reversed in the receiver case. This ordering agrees with the exposure profile ordering in the swap example before.

The fact that the EEs appear slightly decreasing, or at the very least constant (except for the LGM EE in the receiver swaption case) seems to be inconsistent with the increasing PFEs for both payer and receiver contracts. But this can again be explained in terms of the change of drift involved in changing measure. As one moves from the MMA measure to the 20-Forward and 60-Forward measures, the distribution of the short rate at each time horizon shifts down the real line as a result of the lower drift of the short rate under these two measures. Consequently the payer swaption, the value of which is positively related to the short rate, gets pulled more out-the-money under the 20-Forward and 60-Forward measures than under the MMA measure, and thus we see decreasing EE profiles for the payer swaption under these two forward measures. This effect is most pronounced under the 60-Forward measure, as the short rate drift is smaller under this measure than under the MMA and 20-Forward measures respectively. The PFEs, meanwhile, are increasing since they only consider the 95th percentile of the short rate paths under the different measures which, despite the reduction in drift as one moves to the forward measures, is still increasing. That is to say that even though the average of the swaption values at each time horizon tends to be constant, or perhaps decreasing, there are still swaption value paths in the 95th percentile of paths which increase over time. Likewise, we see a reversal of the ordering of the EEs for the receiver swaption, as its value is negatively related to the short rate, and thus has a lower EE for higher drifts. The PFEs for the receiver swaption under the MMA, 20-Forward and 60-Forward measures are also increasing, but attain much lower levels than those of the payer swaption as a result of the positively trending short rate.

Chapter 5

Least Squares Monte Carlo

5.1 The Modified LSM Algorithm

We present here a modification of the LSM algorithm originally due to Longstaff and Schwartz (2001) which will be used in order to estimate contract exposures. This is based on the algorithm given in a paper by Schöftner (2008), but it is extended to include intermediate cashflows.

Suppose we have a contract with value $V(t, X_t)$, where X_t is the underlying risk factor, with payoff at maturity T of $V(T, X_T) = h(X_T)$. Further suppose that the contract pays discrete cashflows $C(t, X_t)$ at times $0 < \tau_1 < \tau_2 < ... < \tau_k \leq T$. The value of this contract at some time $t < T$ is given by

$$V(t, X_t) = \mathrm{E}_{\mathbb{Q}^N}\left[V(s, X_s)\frac{N(t, X_t)}{N(s, X_s)} + \sum_{t<\tau_i\leq s} C(\tau_i, X_{\tau_i})\frac{N(t, X_t)}{N(\tau_i, X_{\tau_i})} \middle| \mathcal{F}_t \right], \qquad (5.1)$$

where $\mathbb{Q}^N$ is the EMM for the numéraire $N(t, X_t)$, $t \leq s \leq T$, and $C(s, X_s) = 0$ if $s \notin \{\tau_1, ..., \tau_k\}$. If this conditional expectation cannot be calculated in closed-form, then one would have to resort to calculating it by Monte Carlo simulation. This becomes especially computationally intensive if one wishes to compute multiple realisations of the contract value at a set of future time horizons. The LSM algorithm, however, avoids this issue by instead estimating this conditional expectation through cross-sectional regression against the information available at time t. Thus we have that

$$\mathrm{E}_{\mathbb{Q}^N}\left[V(s, X_s)\frac{N(t, X_t)}{N(s, X_s)} + \sum_{t<\tau_i\leq s} C(\tau_i, X_{\tau_i})\frac{N(t, X_t)}{N(\tau_i, X_{\tau_i})} \middle| \mathcal{F}_t \right] \approx \sum_{j=1}^{m} \beta_j \phi_j(X_t), \qquad (5.2)$$

where $\beta_1, ..., \beta_m$ are the regression coefficients at time t, and $\phi_1, ..., \phi_j$ a set of suitable regression basis functions.

In approximating a contract's value over time, it would be advantageous not to have to compute the sum of future cashflows which appears in (5.1). We can ensure that we need only consider at most one cashflow at each simulation horizon by requiring that the cashflow times, $\{\tau_1, ..., \tau_k\}$, be a subset of the simulation time grid, $\{0 = t_0, t_1, t_2, ..., t_n = T\}$. If this requirement is met, then we can simulate L realisations of a contract's value under $\mathbb{Q}^N$ as follows:

1. Simulate L independent paths $X^{(l)}_{t_i}, i = 0, ..., n\ ;\ l = 1, ..., L$ under $\mathbb{Q}^N$.
2. Compute the value, $N(t_i, X^{(l)}_{t_i})$, of the numéraire along each path.
3. Set $\hat{V}(t_n, X^{(l)}_{t_n}) = h(X^{(l)}_{t_n})$ for $l = 1, ..., L$.
4. For $i = n-1, n-2, ..., 1$ set $\hat{V}(t_i, X^{(l)}_{t_i}) = \sum_{j=1}^{m} \beta_j \phi_j(X^{(l)}_{t_i})$, where $\beta_1, ..., \beta_m$ are obtained by regressing

$$\left(\hat{V}(t_{i+1}, X^{(l)}_{t_{i+1}}) + C(t_{i+1}, X^{(l)}_{t_{i+1}})\right) \frac{N(t_i, X^{(l)}_{t_i})}{N(t_{i+1}, X^{(l)}_{t_{i+1}})}$$

   against the set of regression basis functions, $\phi_1, ..., \phi_j$ evaluated at $X^{(l)}_{t_i}, l = 1, ..., L$. Note that $C(t_{i+1}, X^{(l)}_{t_{i+1}}) = 0$ if $t_{i+1} \notin \{\tau_1, ..., \tau_k\}$

The result of applying this algorithm is a set of paths, $\hat{V}^{(l)}_{t_i}$, which can then be used to calculate the exposure measures of interest.

In the analysis that follows, we use the same discretised time grids, term structure of instantaneous forward rates (equivalently the term structure of zero-coupon bond prices), and model parameters as before. For the regression basis we choose the first three Laguerre polynomials given by

$$\phi_n(x) = \frac{e^x}{n!} \frac{d^n}{dx^n} \left(e^{-x} x^n\right) , \tag{5.3}$$

for $n = 1, 2, 3$. This choice of basis functions is motivated by Longstaff and Schwartz (2001), who use the same set of functions with an additional weighting factor of $e^{-x/2}$.

5.2 Measuring Goodness of Fit

Naturally, given a set of estimated contract paths, one would want to measure how accurately these paths approximate the actual paths. To this end we make use of the goodness of fit measures defined in Schöftner (2008).

Suppose we have L simulated risk-factor paths at times $t_1, t_2, ...t_n$. Denote the true value of a contract on path l at t_i by $V^{(l)}_{t_i}$, its estimate by $\hat{V}^{(l)}_{t_i}$, and the cross-sectional mean at t_i by $\bar{V}_{t_i}$. We can then define the following goodness of fit measures at time t_i:

Mean Squared Error (MSE):

$$\mathrm{MSE}_{t_i} = \frac{1}{L}\sum_{k=1}^{L}(V_{t_i}^{(k)} - \hat{V}_{t_i}^{(k)})^2 \,. \tag{5.4}$$

Relative Squared Error (RSE):

$$\mathrm{RSE}_{t_i} = \frac{\sum_{k=1}^{L}(V_{t_i}^{(k)} - \hat{V}_{t_i}^{(k)})^2}{\sum_{k=1}^{L}(V_{t_i}^{(k)} - \bar{V}_{t_i})^2} \,. \tag{5.5}$$

Whereas MSE measures the average squared deviation in terms of the units in which V_t is expressed, the RSE divides the aggregate squared deviations by the sample variance of V_t. Thus the RSE is, in a sense, scale-free and should allow us to compare goodness of fit for the LSM approximations under different probability measures.

We can similarly define the Average Mean Squared Error (AvMSE) and Average Relative Squared Error (AvRSE):

$$\mathrm{AvMSE} = \frac{1}{n}\sum_{i=1}^{n}\mathrm{MSE}_{t_i} \quad \text{and} \quad \mathrm{AvRSE} = \frac{1}{n}\sum_{i=1}^{n}\mathrm{RSE}_{t_i} \,,$$

which each provide a single summary measure of the goodness of fit across all paths and time horizons.

5.3 Swap Exposure

On the whole the LSM algorithm presented before seems, in most cases, to provide a satisfactory approximation of the swap exposure profiles, as can be seen in Figure 5.1 and Figure 5.2. The accuracy of these approximations, however, seems to depend on the contract type (i.e. payer or receiver), the chosen probability measure, and the degree of polynomial regression basis employed in the LSM algorithm. For the sake of convenience, in the discussion that follows the phrase "n-th degree LSM exposure profile" is taken to mean "the LSM exposure profile approximation obtained when using an n-th degree Laguerre polynomial regression basis".

The first thing to note is that there is no visible difference between second-degree and third-degree LSM exposure profile approximations, with the green line (which corresponds to the third-degree polynomial basis) lying along the blue line (corresponding to the second-degree polynomial basis) at all points. For the payer swap, the approximated EE profiles are almost indistinguishable from the actual EE profile, across all four measures (Figure 5.1). The PFE under the MMA and 60-Forward measures seems best approximated, at all time horizons, by a first-degree

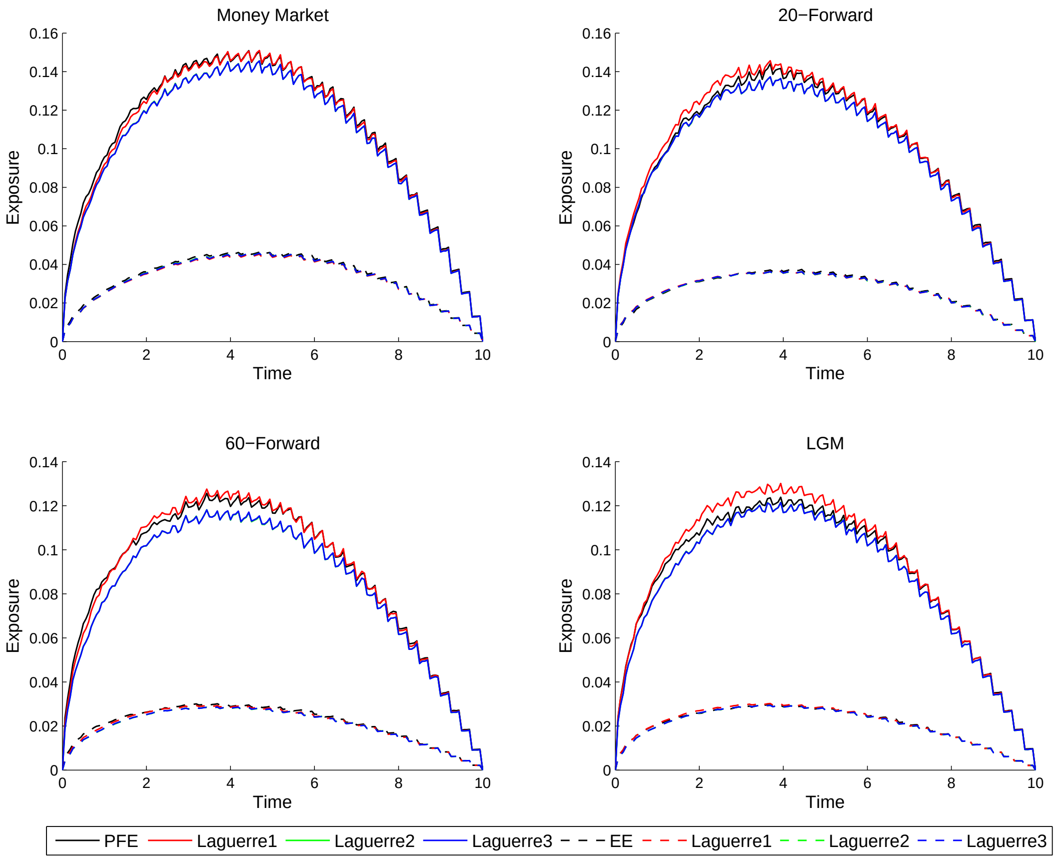

Fig. 5.1: LSM approximations of 95% PFE and EE for a 10 year vanilla payer swap.

LSM PFE, whereas the 20-Forward and LGM PFEs seem to switch between being well approximated by a first-degree LSM PFE and a second-degree or third-degree LSM PFE. The maximum relative error, across all four measures, between the actual PFE profiles and the best approximating PFE profiles is only about 10% in the payer case.

The approximations fare slightly worse in the receiver case (Figure 5.2), with there being visible differences in EEs and quite marked differences in PFEs (by about 20%, in relative terms). Here the higher degree LSM approximations (i.e. second-degree and third-degree approximations) appear to provide a better fit of the PFEs than the first-degree LSM approximations. In fact there seems to be a positive bias in the first-degree LSM approximation, relative to higher order polynomial approximations, since the first-degree LSM profiles lie above the higher degree profiles in the payer case, and below the higher degree profiles in the receiver case.

What seems strange is that the LSM PFE approximations for the receiver swap

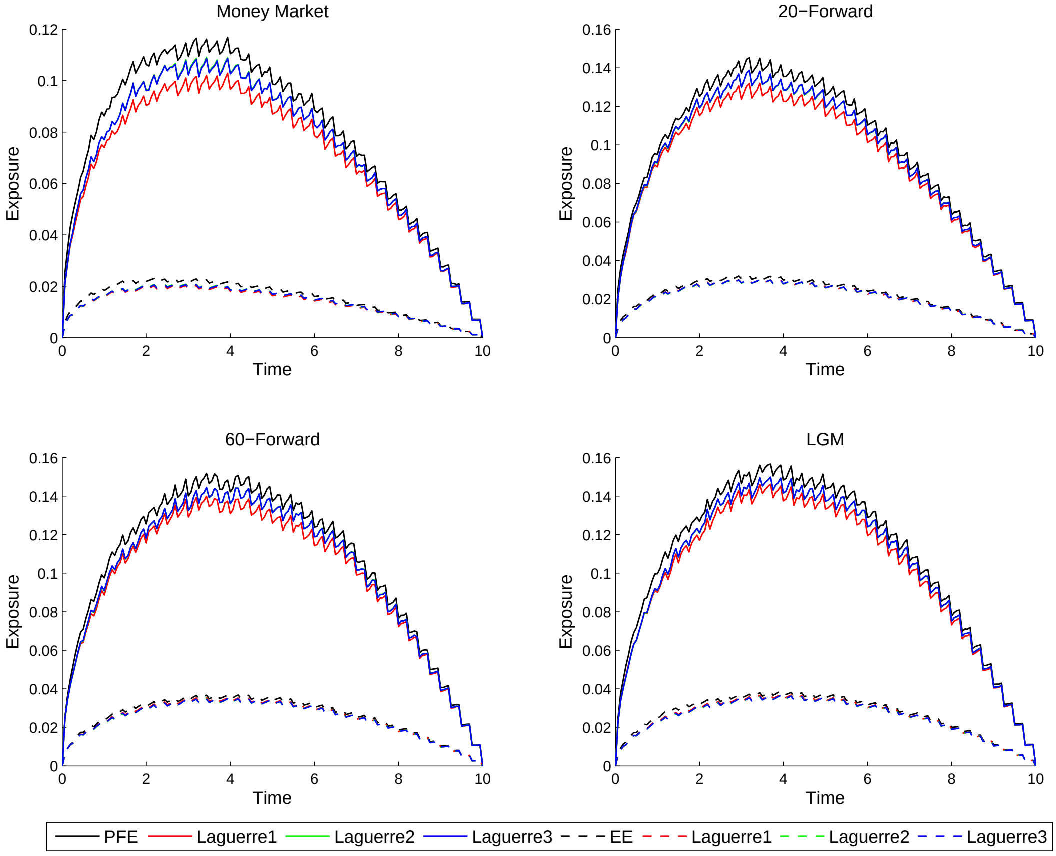

Fig. 5.2: LSM approximations of 95% PFE and EE for a 10 year vanilla receiver swap.

are much worse than the payer swap under the MMA measure, whereas the accuracy does not deteriorate as much under the 20-Forward and 60-Forward measures. This, however, can be explained by the relative drifts of the short rate process under these measures. Because the short rate has a higher drift under the MMA measure, relative to the other measures, there tends to be more positive realised payer swap values at each time horizon than negative ones (which are, equivalently, positive receiver swap values). Thus the regression run at each stage of the LSM process would be biased towards the positive payer swap values, and as a result provide a potentially worse fit of the negative payer (positive receiver) swap values. Moreover, these errors would accumulate and compound as the LSM algorithm iterates backwards through time. Since the short rate has a lower drift under the 20-Forward and 60-Forward measures, this problem is not as pronounced (as there are relatively fewer positive payer swap values in the simulations under these measures) and, as a result, the approximations of the payer swap PFEs are better than those obtained under the MMA measure.

Rather than relying solely on visual inspection, we can also quantify the goodness of fit of the LSM approximations by calculating the AvMSE and AvRSE for the LSM approximated price paths under each of the chosen probability measures. The AvMSE and AvRSE for the swap price paths are given in Table 5.1. On both an absolute and relative basis the price paths estimated using a second-degree polynomial basis provide the best fit under all four probability measures. Comparing relative fits, we see that the error in approximating price paths is minimised when working under the 20-Forward measure. These values seem at odds with the apparent goodness of fit of the LSM swap EE and PFE approximations, especially when comparing the LSM PFEs to the actual PFEs in the payer case, where the first-degree polynomial regression basis seems to yield a better fit. This could potentially be as a result of the higher-degree polynomial regressors overfitting the contract values, especially at the tail ends of the regression data (which are precisely the values we are concerned with in approximating PFEs).

Tab. 5.1: LSM swap goodness of fit statistics.

	AvMSE				AvRSE			
	MMA	20-Forward	60-Forward	LGM	MMA	20-Forward	60-Forward	LGM
Laguerre1	2.5165E-05	1.8738E-05	1.7671E-05	1.7906E-05	0.0067	0.0047	0.0045	0.0044
Laguerre2	**1.2888E-05**	**7.9483E-06**	**1.3475E-05**	**1.0299E-05**	**0.0042**	**0.0027**	**0.0043**	**0.0033**
Laguerre3	1.3159E-05	8.3986E-06	1.3662E-05	1.0542E-05	0.0042	0.0028	0.0044	0.0033

5.4 Swaption Exposure

The LSM approximations of the swaption exposures seem also to be quite satisfactory, although the accuracy does again depend on the polynomial degree, the contract type, and the probability measure (with the LSM exposure profile approximations faring far worse under the LGM measure than the approximations under other measures).

For both payer (Figure 5.3) and receiver (Figure 5.4) swaptions, the first-degree LSM exposure approximations fare significantly worse than the second and third-degree approximations. This is not unexpected, however, as the value of a swaption is a non-linear function of its underlying risk factor (e.g. the short rate, or the LGM Gaussian state process X_t). Thus, at the regression stage of the LSM algorithm, the first-degree polynomial basis results in a worse estimation of the contract value than the higher degree bases, which results in an overall poorer approximation of the contract value paths. The second and third-degree LSM exposure approxima-

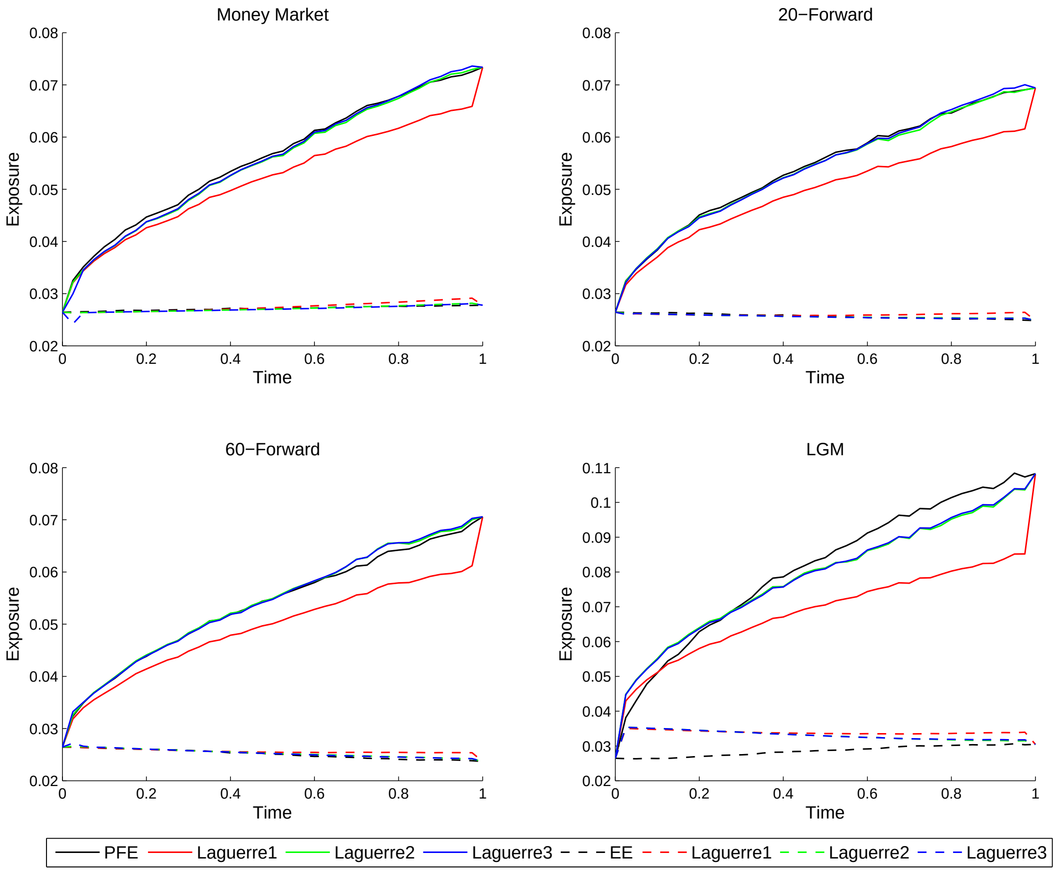

Fig. 5.3: LSM approximations of 95% PFE and EE for a 1Y5Y year payer swaption.

tions are very similar in the payer case, but have slightly more pronounced differences in the receiver case. As a result of the poor performance of the first degree LSM approximation, we will disregard it in the discussion that follows.

On the whole the payer swaption exposures seem better approximated by the LSM algorithm than the receiver swaption exposures (the LGM case not withstanding). The maximum relative error in the PFEs of the payer swaption under the MMA, 20-Forward and 60-Forward measures is only 2.5%, whereas this error grows to about 20% for the receiver swaption PFEs (although the PFE approximations do improve as we move backwards in time from the contracts' maturities). We can attribute this difference in fit to the moneyness of the swaption values being simulated, and the different drifts of the short rate under the four probability measures. The short rate has a positive drift under these measures, which means that, on average, it takes on higher values in the future than its initial value, and hence the simulated payer swaptions tend to spend more time in-the-money than the receiver swaptions. This results in a poorer estimation of the receiver swaption

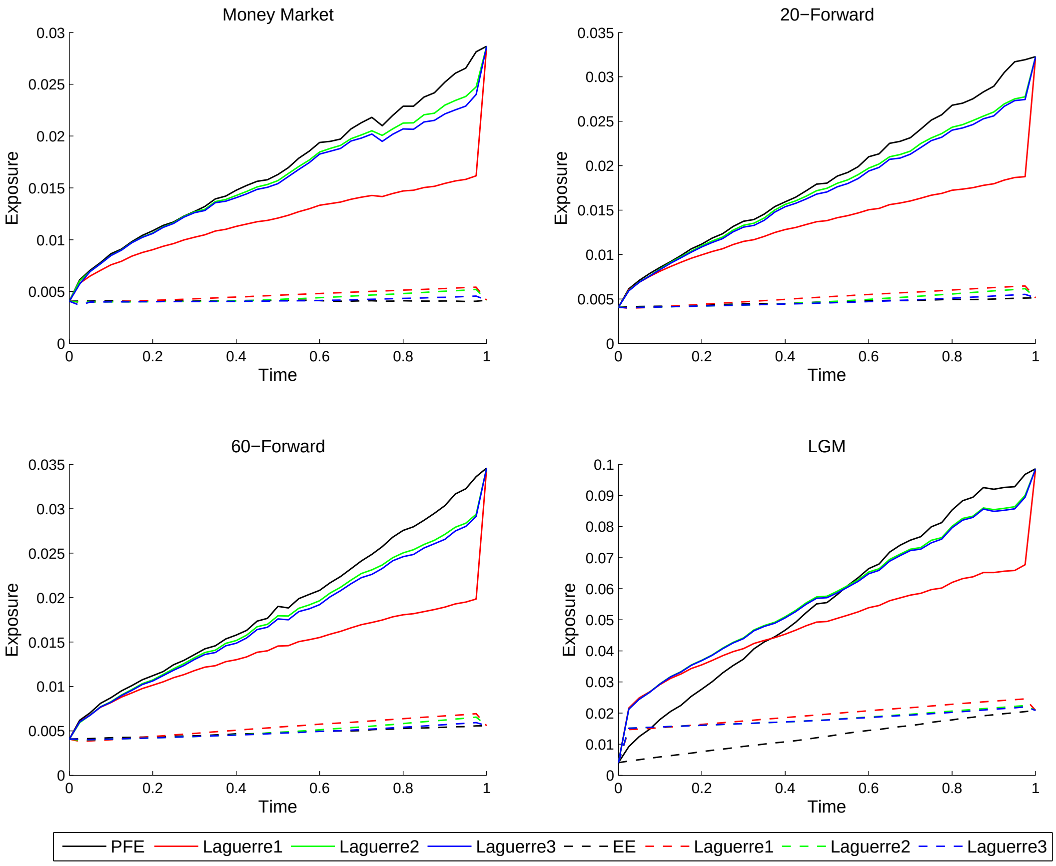

Fig. 5.4: LSM approximations of 95% PFE and EE for a 1Y5Y year receiver swaption.

values relative to the payer swaption values, and hence the worsened approximation of their respective exposure profiles. Moreover, as we move from the MMA to the two forward measures, the drift of the short rate is reduced, and hence, in theory, the exposure approximations should worsen for the payer swaption (as now fewer paths are in-the-money) and improve for the receiver swaption. In practice, the reduction in the drift as we move from the MMA to the 20-Forward and 60-Forward measures is small, and thus this change is imperceptible. The importance sampling method could be used to improve the accuracy of the exposure approximations by reducing the number of out-the-money paths which are simulated.

The reasons for the relatively poor approximations of the swaption exposures under the LGM measure are not immediately apparent. We note that the approximated PFEs for both payer and receiver swaption initially underestimate the true PFE, but cross over and end up over estimating the PFEs as we move backwards in time from maturity. Likewise the EE profile approximations also exhibit an upwards bias as we move backwards in time. Thus it is possible that the LSM algorithm introduces an upwards bias in approximated swaption exposures under the LGM measure, but further investigation would be required to determine how, exactly, this bias is introduced.

Tab. 5.2: LSM swaption goodness of fit statistics.

		AvMSE				AvRSE			
		MMA	20-Forward	60-Forward	LGM	MMA	20-Forward	60-Forward	LGM
Payer	Laguerre1	1.4429E-05	1.5912E-05	1.6415E-05	1.4553E-04	0.0372	0.0429	0.0460	0.2587
	Laguerre2	2.9347E-06	2.9961E-06	2.6267E-06	4.8548E-05	**0.0072**	0.0076	0.0066	0.1623
	Laguerre3	**8.7753E-07**	**8.7858E-07**	**8.5432E-07**	**4.5496E-05**	0.0125	**0.0027**	**0.0033**	**0.1599**
Receiver	Laguerre1	1.7829E-05	1.9415E-05	2.0016E-05	1.9476E-04	0.3049	0.2897	0.2874	1.3494
	Laguerre2	2.8670E-06	2.9629E-06	2.5978E-06	5.4912E-05	0.0406	0.0368	0.0314	1.1129
	Laguerre3	**6.5749E-07**	**8.0897E-07**	**8.2753E-07**	**5.1750E-05**	**0.0117**	**0.0095**	**0.0100**	**1.1089**

If we consider the goodness of fit statistics of the swaption price path approximations we see that on both a relative basis (except in the MMA case), and on an absolute basis the third-degree LSM approximations result in the best overall fit (Table 5.2). As expected, the first-degree approximations fare worse than the second and third-degree approximations due to the non-linear relationship between the value of a swaption and its underlying risk factor. The AvRSE values highlight the poorer overall approximation of the receiver swaption paths compared to the payer swaption paths which, as has been mentioned, is due to the relatively greater number of out-the-money receiver swaption paths. Moreover, the significantly inflated AvRSE values under the LGM measure indicate that swaption value paths are poorly approximated under this measure. This agrees with the poor exposure profile approximations that were obtained earlier.

5.5 Introducing a Change of Measure

Whilst the modified LSM algorithm presented before allows one to avoid having to run Monte Carlo simulations within Monte Carlo simulations to compute exposure profiles for contracts without closed-form valuation formulae, it still does not address the fundamental problem raised in the preceding section: namely that one can change the shape of an exposure profile by simply working under a different EMM (Stein, 2014).

As has been mentioned, for risk management purposes we need to compute exposures under the real-world measure. However, the modified LSM algorithm cannot be directly applied to real-world risk factor simulations as it based on the principal of martingale valuation, which is achievable only under an EMM. If we, however, knew the change of measure required to move from the real-world measure to an EMM, we could adapt the previously presented LSM algorithm in the manner outlined by Schöftner (2008):

1. Simulate L independent paths $X_{t_i}^{(l)}, i = 0, ..., n, l = 1, ..., L$ under $\mathbb{P}$ (not $\mathbb{Q}^N$) and compute the value, $N(t_i, X_{t_i}^{(l)})$, of the numéraire along each path. In addition compute the values, under $\mathbb{P}$, of any intermediate cashflows $C(t_i, X_{t_i}^{(l)})$.

2. Compute the value of the time-discretised Radon-Nikodym process

$$R(t_i, X_{t_i}^{(l)}) = \exp\left(\sum_1^i \lambda_{t_{i-1}}(t_i - t_{i-1})Z_i^{(l)} - \frac{1}{2}\sum_1^i \lambda_{t_{i-1}}^2(t_i - t_{i-1})\right),$$

where λ_t is the kernel of the Girsanov transformation which effects the change of measure from $\mathbb{P}$ to $\mathbb{Q}^N$, and the $Z_i^{(l)}$ are the independent standard normal variables used to simulate the l-th risk factor path.

3. In the regression step, we now regress

$$\left(\hat{V}(t_{i+1}, X_{t_{i+1}}^{(l)}) + C(t_{i+1}, X_{t_{i+1}}^{(l)})\right) \frac{N(t_i, X_{t_i}^{(l)})}{N(t_{i+1}, X_{t_{i+1}}^{(l)})} \frac{R(t_{i+1}, X_{t_{i+1}}^{(l)})}{R(t_i, X_{t_i}^{(l)})}$$

against the regression basis functions evaluated at $X_{t_i}^{(l)}$.

This algorithm estimates the value of the contract at time t through cross-sectional regression as

$$\begin{aligned} V(t, X_t) &= \mathrm{E}_{\mathbb{Q}^N}\left[V(s, X_s)\frac{N(t, X_t)}{N(s, X_s)} + \sum_{t<\tau_i\le s} C(\tau_i, X_{\tau_i})\frac{N(t, X_t)}{N(\tau_i, X_{\tau_i})} \middle| \mathcal{F}_t \right] \\ &= \mathrm{E}_{\mathbb{P}}\left[V(s, X_s)\frac{N(t, X_t)}{N(s, X_s)}\frac{R(s, X_s)}{R(t, X_t)} + \sum_{t<\tau_i\le s} C(\tau_i, X_{\tau_i})\frac{N(t, X_t)}{N(\tau_i, X_{\tau_i})}\frac{R(\tau_i, X_{\tau_i})}{R(t, X_t)} \middle| \mathcal{F}_t \right] \\ &\approx \sum_{j=1}^m \beta_j \phi_j(X_t)\,, \end{aligned} \tag{5.6}$$

where $t \le s \le T$, and the β_j and ϕ_j are the regression coefficients and regression basis functions respectively. Note that we require the change of measure function, $R(t, X_t)$, because we are now considering realisations of X_t under $\mathbb{P}$, not $\mathbb{Q}^N$. Moreover, since regression is done against realisations of X_t under $\mathbb{P}$, we now obtain contract value paths under $\mathbb{P}$.

Moving from the real-world measure to the risk-neutral measure requires estimating the market price of risk, as this is the Girsanov kernel required to effect the change of measure. In practise, there may be difficulties in estimating the market price of risk, which would limit the effectiveness of the LSM algorithm. However, we do not address the issue of estimating the market price of risk in this dissertation.

In order for us to illustrate the change of measure process, we assume that the risk neutral measure coincides with the real world measure (i.e. that the real world dynamics of the short rate are precisely the dynamics of the short rate under the MMA measure), and we fix the 20-Forward measure as our chosen EMM under which contracts will be valued. As a result, the required Girsanov kernel is simply $-\sigma B(t, 20)$, which is the volatility of the 20 year zero-coupon bond. We follow the same methodology as the previous section in order to compute swap and swaption exposures, but now incorporate the changes to the LSM algorithm given above. For both contract types we simulate the short rate under the MMA and 20-Forward measures, as well as the time-discretised Radon-Nikodym derivative, using the same set of standard normal realisations. In addition, the intermediate cashflows for the swaps are computed from the zero-coupon bond curves associated with each realisation of the short rate under the MMA measure.

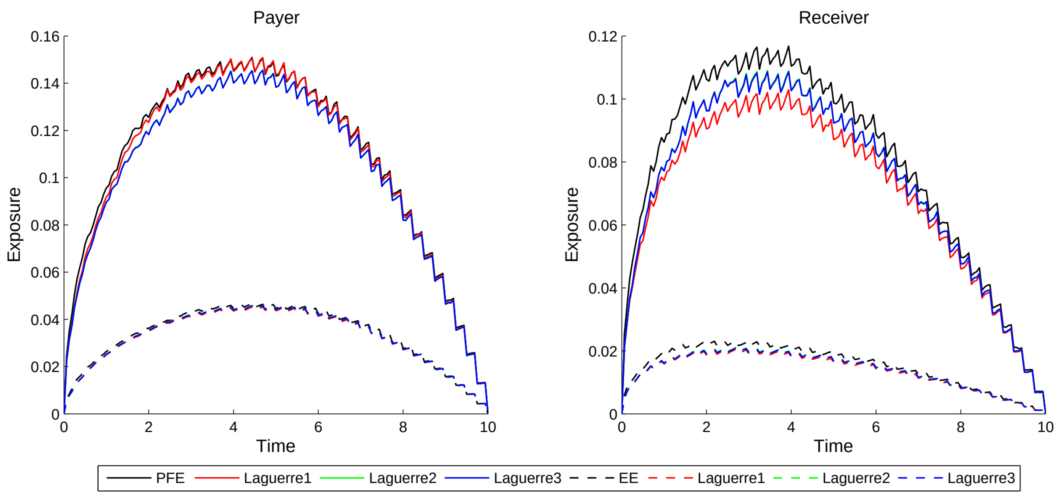

Fig. 5.5: LSM approximations of 95% PFE and EE for a 10 year vanilla payer and receiver swaps (MMA measure via 20-Forward measure).

For both the payer and the receiver swaps (Figure 5.5), the “real-world measure” exposure profiles obtained via a change of measure from the 20-Forward measure to the MMA measure are, on visual inspection, almost the same as those obtained directly under the MMA measure. This can be confirmed by looking at the relative errors of both the PFEs and EEs obtained (Figure 5.6 and Figure 5.7), where the relative errors when using a change of measure coincide with the direct simulation errors. The payer PFE and EE relative errors are greatest just after the contract's initiation, at about 11% and 12% respectively. The relative errors in the receiver case grow to maximums of 17% and 19% for the PFE and EE respectively. The similarity of the exposure profiles obtained via a change of measure to those obtained directly under the MMA measure is most likely due to the fact that at each cashflow date the cashflow paid on that date must be included in the LSM calculation, with these cashflows being computed directly under the MMA measure. This seems to offset, at least partially, the errors arising as a result of the change of measure.

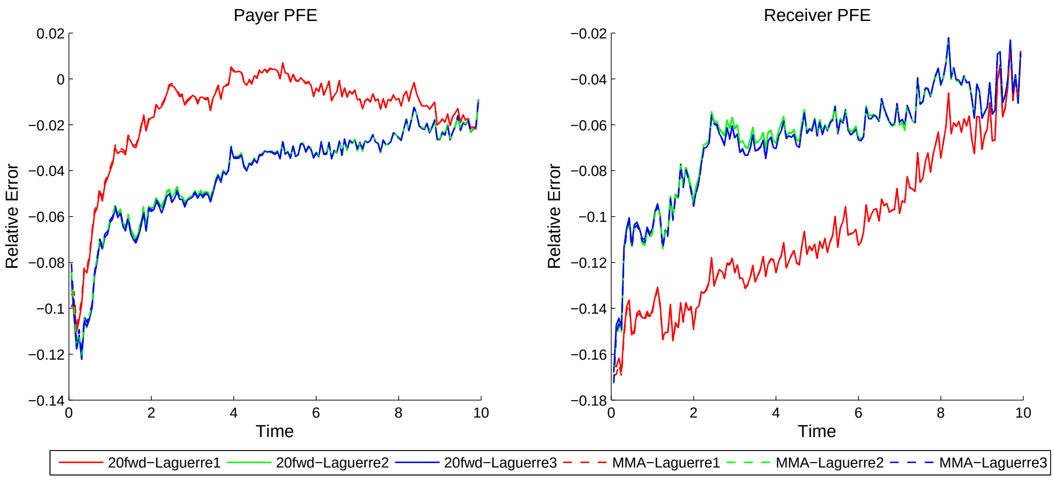

Fig. 5.6: Relative errors of LSM approximations of 95% PFE for 10 year vanilla payer and receiver swaps (MMA measure via 20-Forward measure).

In the case of the swaptions, we see again a worse approximation of the exposure profiles in the receiver case than in the payer case (Figure 5.8). This is again a result of there being relatively more in-the-money payer swaption paths than in-the-money receiver swaption paths. If anything, the approximations of the receiver swaption exposures would be worse when the change of measure is introduced, as the higher drift of the short rate process under the MMA measure results in fewer in-the-money paths to regress against.

As one can see from the plots of relative errors (Figure 5.9 and Figure 5.10), the approximations of the exposure profiles under the MMA measure, when working with a change of measure, are worse than those of the LSM approximations computed entirely under the MMA measure. Moreover, the exposure profile approximations worsen as we move backwards in time from the contract maturity date.

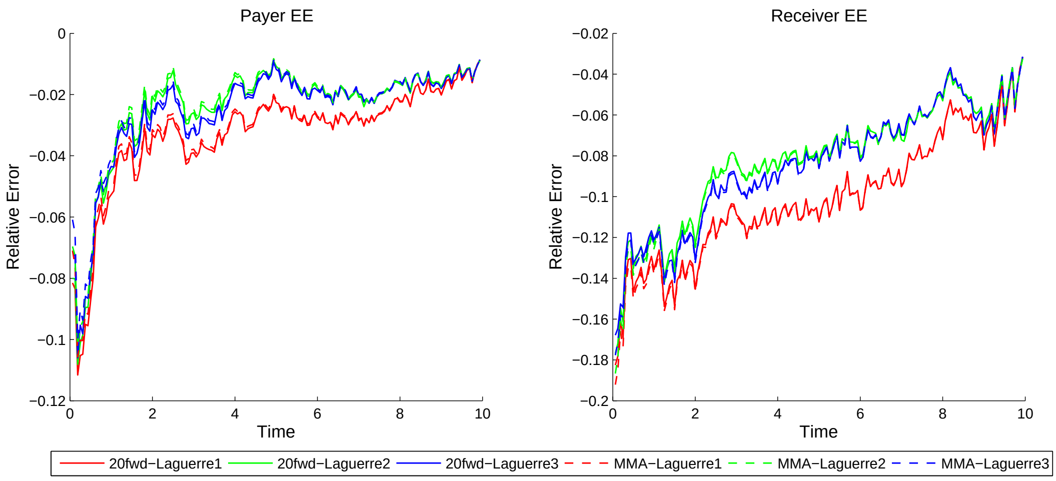

Fig. 5.7: Relative errors of LSM approximations of EE for 10 year vanilla payer and receiver swaps (MMA measure via 20-Forward measure).

This is presumably as a result of the approximate nature of the time-discretised Radon-Nikodym derivative used to effect the change of measure, as well as the accumulation of these errors as the LSM algorithm iterates backward through time. Ignoring, once again, the first-degree LSM approximation, we see that for the payer swaption, both the PFE and EE relative errors under the change of measure remain within 10% for the vast majority of the life of the contract (in fact the relative errors peak at about 12) . The errors in the receiver case are far more pronounced, due to the interplay of the approximate change of measure function, as well as the moneyness issue mentioned before. The errors remain within 10% for only about a third of the life of the contract for the EE, and about half of the life of the contract for the PFE. The EE error reaches a maximum of about 30%, whereas the PFE error increases to a maximum of about 20%. It is worth noting that the largest relative errors for the approximate PFEs occur early in the lives of the contracts, where the PFEs are actually at their lowest, instead of closer to the maturities of the contracts, where PFEs are at their highest.

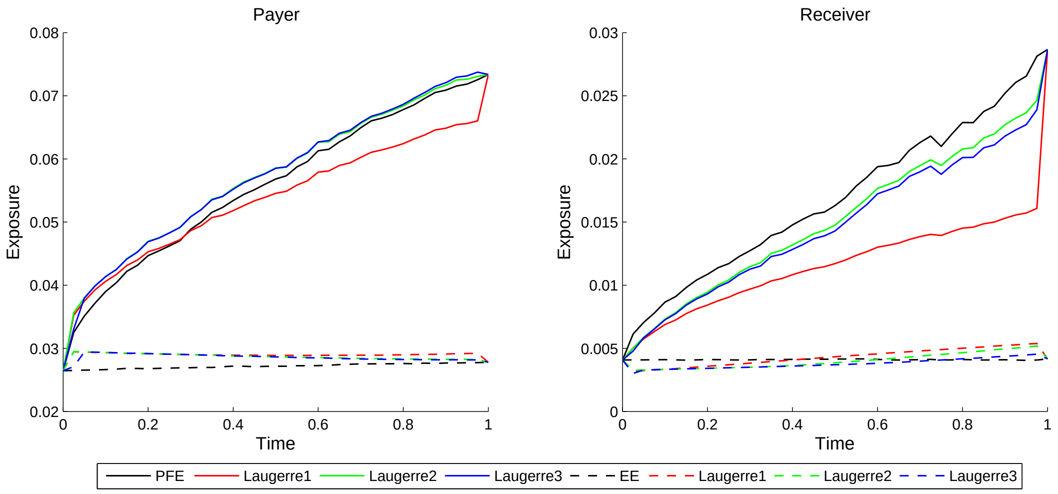

Fig. 5.8: LSM approximations of 95% PFE and EE for 1Y5Y year payer and receiver swaptions (MMA measure via 20-Forward measure).

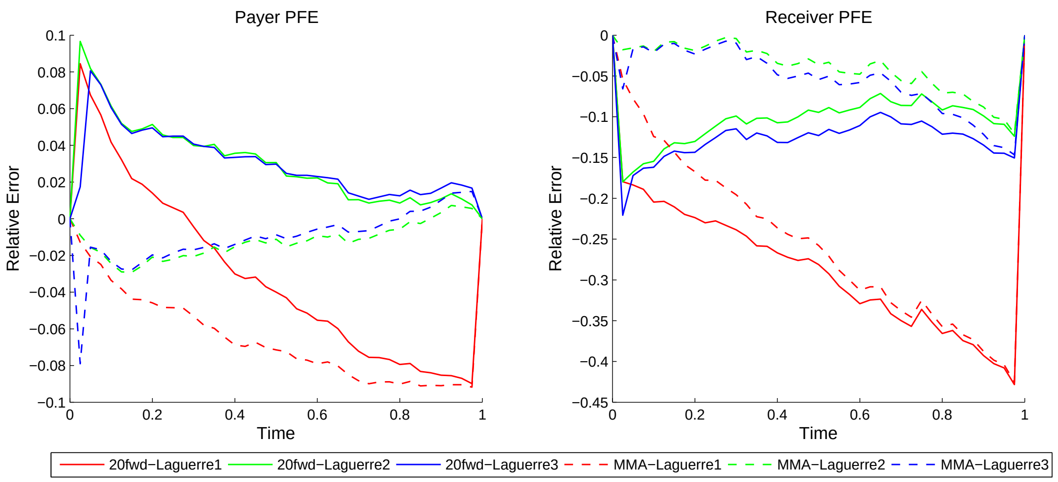

Fig. 5.9: Relative errors of LSM approximations of 95% PFE for 1Y5Y year payer and receiver swaptions (MMA measure via 20-Forward measure).

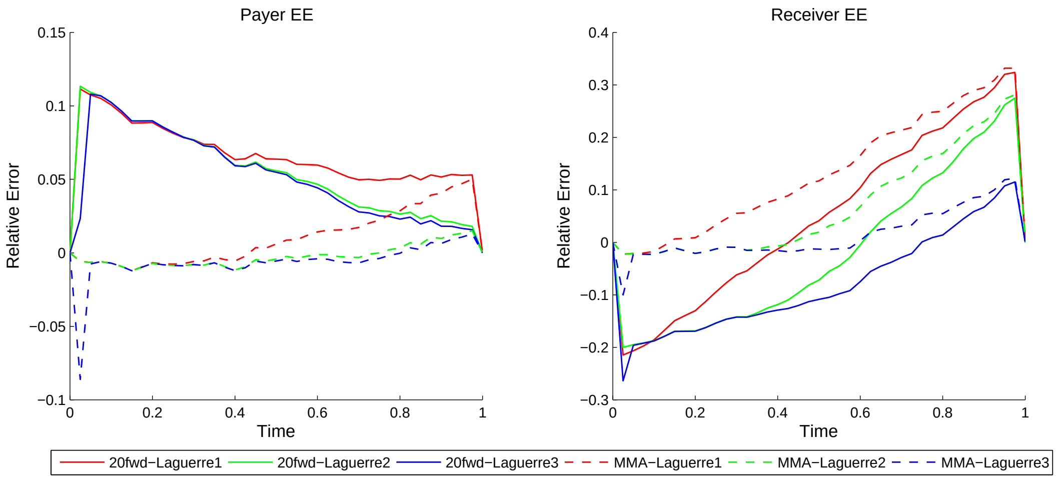

Fig. 5.10: Relative errors of LSM approximations of EE for 1Y5Y year payer and receiver swaptions (MMA measure via 20-Forward measure).

5.6 Comparison of Computation Times

Table 5.3 summarises the time taken to compute contract price paths. Here "Closed-Form" refers to using closed-form valuation formulae to value the contracts at each time horizon, whereas "LSM" refers to the LSM approximated price paths and "LSM Change of Measure" refers to the price paths obtained via a change of measure.

Tab. 5.3: Closed-form and LSM price path computation times

		Runtime (seconds)			
		MMA	20-Forward	60-Forward	LGM
Closed-Form	Swap	797.64	800.95	799.81	1201.66
	Payer Swaption	2234.23	2229.32	2205.84	812.38
	Receiver Swaption	2231.34	2229.25	2209.58	813.7
LSM	Swap	105.12	511.345	508.95	189.28
	Payer Swaption	56.01	159.27	158.83	21.62
	Receiver Swaption	55.58	160.11	159.08	21.86
LSM Change of Measure	Swap	515.17	-	-	-
	Payer Swaption	156.56	-	-	-
	Receiver Swaption	157.16	-	-	-

The first thing to note about these run times is that using the LSM algorithm reduces the time taken to produce swap price paths by a factor of 1.5 to 7.5 (depending on the choice of numéraire) and reduces the time taken to compute swaption price paths by a factor of 13 to 40 times. The significantly higher computation times for the closed-form algorithm are due to the increased number of interpolations of the term structure of zero-coupon bond prices at each time horizon, with this being required to value both swaps and swaptions. In addition, the swaption valuation is slowed down further as a result of the root finding method used to determine the "critical value" in Jamshidian's trick. The intermediate step of zero-coupon bond option pricing which is required to price swaptions under the HW model adds additional computation time under the MMA, 20-Forward and 60-Forward measures.

Within the LSM approximations, computation time depends on whether any intermediate cashflows need to be computed (as is the case for swaps) and the chosen numéraire (since some numéraires, such as the T-forward numéraire, are more costly to compute than others). It must be noted that the bulk of the LSM computation time is due to these two processes, as the total time to perform all the regressions takes, on average, only 1-3 seconds.

The swap LSM price paths which are obtained via a change of measure show a moderate speed increase of 1.5 times relative to the closed-form paths, whereas the swaption price path generation is sped up by as much as 14 times. The increase in the computation times relative to the LSM paths approximated directly under the MMA measure is a result of having to compute the 20-Forward numéraire paths, and also having to compute the time-discretised Radon-Nikodym derivative. As a result, these runtimes are similar to those of the LSM algorithm under the 20-Forward measure.

Although there should be no reason to resort to the LSM algorithm for estimating price paths when closed-form valuation formulae are available (since the primary benefit of the LSM algorithm is to avoid the nested Monte Carlo simulations that would arise in the absence of closed-form valuation formulae), one might still choose to employ the LSM algorithm, instead of a closed-form approach, so as to reduce computation time. However it should not be assumed that, as a rule, the LSM price path algorithm will always run faster than the closed-form one based on the results given in this section. Rather, its suitability should be determined on a case-by-case basis.

Chapter 6

Conclusion

In this dissertation we have considered the modelling of PFEs and EEs for vanilla interest rates and swaptions under a selection of commonly used EMMs and their associated numéraires. We have shown that the magnitude and shape of these profiles can vary, sometimes quite substantially, under different probability measures. In the case of a short rate modelled under the HW model, these differences arise from the drift adjustment that is made as we change between the measures. The more negative the drift adjustment, the smaller (greater) the value of instruments paying (receiving) the fixed rates. This applies both to swaps, and their swaption variants. As a result of this, the capital requirements of a firm may differ depending on the probability measure chosen to model exposures.

Despite this, modelling of exposures under an EMM can be beneficial, especially for contracts and portfolios which cannot be valued in closed-form. This is because one can leverage Least Squares Monte Carlo techniques in order to approximate contract values, and thus avoid having to run additional Monte Carlo simulations in order to price contracts. We have described a modified LSM algorithm which can be used to approximate contract exposures. The swap and swaption exposures were adequately approximated by the LSM algorithm, and these approximations turned out to be quicker to compute than the actual exposure profiles. However, we have observed that the accuracy of these approximations can, to a greater or lesser extent, depend on the contract being modelled, the degree of the regression basis used and the measure under which the exposures are computed.

Finally, we have shown how to introduce a change of measure to the LSM algorithm which allows for the approximation of real-world exposure profiles, rather than risk-neutral exposure profiles. This is beneficial not only because of the potential computational savings, but also because these profiles can be used for risk management and for determining capital requirements (unlike the profiles obtained under an EMM).

Bibliography

Brigo, D. and Mercurio, F. (2007). *Interest Rate Models - Theory and Practice: With Smile, Inflation and Credit*, Springer-Verlag.

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

Geman, H., El Karoui, N. and Rochet, J.-C. (1995). Changes of numéraire, changes of probability measure and option pricing, *Journal of Applied probability* **32**(2): 443–458.

Hagan, P. S. (2009). Evaluating and hedging exotic swap instruments via LGM, *Technical report*, Bloomberg Technical Report.

Hagan, P. S. and Woodward, D. E. (1999). Markov interest rate models, *Applied Mathematical Finance* **6**(4): 233–260.

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

Jamshidian, F. (1989). An exact bond option formula, *The Journal of Finance* **44**(1): 205–209.

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

Pykhtin, M. and Zhu, S. H. (2007). A guide to modeling counterparty credit risk, *GARP Risk Review* (37).

Schöftner, R. (2008). On the estimation of credit exposures using regression-based monte carlo simulation, *The Journal of Credit Risk* **4**(4).

Stein, H. J. (2014). Fixing underexposed snapshots - proper computation of credit exposures under the real world and risk neutral measures, *Available at SSRN: http://ssrn.com/abstract=2365540* .

Appendix A

MATLAB Code

GOF.m

```
function out = GOF(actualPaths,estimatedPaths,prctile,output
    )
%GOF calculates a variety of 'squared-error' goodness of fit
%measures, either across all paths, or only those above a
%specified percentile
%
%   ActualPaths     : realised asset price paths
%   EstimatedPaths : asset price paths estimated from LSM
%                     algorithm
%   prctile         : percentile (%)

[n,N] = size(actualPaths);
if isempty(prctile)
    MSE = sum((actualPaths-estimatedPaths).^2,2)/N;
    AvMSE = sum(MSE)/n;
    RSE = sum((actualPaths-estimatedPaths).^2,2)./sum((
       actualPaths-...
        repmat(mean(actualPaths,2),1,N)).^2,2);
    AvRSE = sum(RSE)/n;
else
    prctileInd = floor(prctile*N);
    [sortedActualPaths, sortedInd] = sort(actualPaths,2);
    MSE = zeros(n,1);
    RSE = zeros(n,1);
    for i = 1:n

        MSE(i) = sum((actualPaths(i,sortedInd(i,prctileInd:
           end))-...
            estimatedPaths(i,sortedInd(i,prctileInd:end)))
               .^2,2)/(N-prctileInd);
        RSE = sum((actualPaths(i,sortedInd(i,prctileInd:end)
           )-estimatedPaths(i,sortedInd(i,prctileInd:end)))
```

```
                 .^2,2)./...
                  sum((actualPaths(i,sortedInd(i,prctileInd:end))-
                      repmat(mean(actualPaths(i,sortedInd(i,
                      prctileInd:end)),2),1,N-prctileInd+1)).^2,2);
        end
        AvMSE = sum(MSE)/n;
        AvRSE = sum(RSE)/n;
end

if strcmpi(output,'mse')
        out = MSE;
elseif strcmpi(output,'AvMSE')
        out = AvMSE;
elseif strcmpi(output,'RSE')
        out = RSE;
elseif strcmpi(output,'AvRSE')
        out = AvRSE;
end

end
```

HullWhiteBondPrices.m

```
function ZCB = HullWhiteBondPrices(t,rt,P0,f0,a,sigma,Tvec)
%HULLWHITEBONDPRICES prices ZCBs under the Hull-White model
%at time t for maturities T1,T2,...TN
%
%   P0      : initial ZCB prices
%   f0      : initial inst. forward rates
%   a       : mean-reversion rate
%   sigma   : short-rate vol
%   t       : settlement
%   Tvec    : maturity dates

Pvec = interp1(P0(:,1),P0(:,2),Tvec);
Pt = interp1(P0(:,1),P0(:,2),t);
ft = interp1(f0(:,1),f0(:,2),t);
B = (1-exp(-a*(Tvec-t)))/a;
A = (Pvec./Pt).*exp(B*ft-(sigma^2/(4*a)).*(1-exp(-2*a*t)).*B
    .^2);
ZCB = A.*exp(-B.*rt);

end
```

HullWhiteRates.m

```
function paths = HullWhiteRates(r0,f0,a,sigma,tVec,N)
%HULLWHITERATES simulates N short rate paths under the
%Hull-White model at times t_1, t_2, ... t_n
%
%   r0       : initial short rate
%   f0       : initial forward curve (f0(1,2)=r0)
%   a        : mean-reversion rate
%   sigma    : vol
%   tvec     : vector of tenors t_0, t_1,...t_n
%   N        : number of paths to simulate

n = length(tVec)-1;
paths = [repmat(r0,1,N); zeros(n,N)];
Z = randn(n,N);

alphaVec = interp1(f0(:,1),f0(:,2),tVec)+(sigma^2)/(2*a^2)
   *(1-exp(-a*tVec)).^2;

for i = 2:n+1
    paths(i,:) =  paths(i-1,:)*exp(-a*(tVec(i)-tVec(i-1)))+
       alphaVec(i)-...
        alphaVec(i-1)*exp(-a*(tVec(i)-tVec(i-1)))+sigma*...
        sqrt((1-exp(-2*a*(tVec(i)-tVec(i-1))))/(2*a))*Z(i
           -1,:);
end

end
```

HullWhiteRatesTfwd.m

```
function paths = HullWhiteRatesTfwd(r0,f0,a,sigma,T,tVec,N)
%HULLWHITERATES simulates N short rate paths in the
%Hull-White model under the T-forward measure at times
%t_1, t_2, ... t_n
%
%   r0       : initial short rate
%   f0       : initial forward curve (f0(1,2)=r0)
%   a        : mean-reversion rate
%   sigma    : vol
%   T        : maturity
%   tvec     : vector of times t_0, t_1,...t_n
%   N        : number of paths to simulate

n = length(tVec)-1;
paths = [repmat(r0,1,N); zeros(n,N)];
Z = randn(n,N);
```

```
alphaVec = interp1(f0(:,1),f0(:,2),tVec)+(sigma^2)/(2*a^2)
   *(1-exp(-a*tVec)).^2;
M = @(s,t) -((sigma/a)^2)*(1-exp(-a*(t-s)))+0.5*((sigma/a)
   ^2)*(exp(-a*(T-t))-...
    exp(-a*(T+t-2*s)));

for i = 2:n+1
    paths(i,:) =  paths(i-1,:)*exp(-a*(tVec(i)-tVec(i-1)))+
       alphaVec(i)-...
        alphaVec(i-1)*exp(-a*(tVec(i)-tVec(i-1)))+M(tVec(i
           -1),tVec(i))+sigma*...
        sqrt((1-exp(-2*a*(tVec(i)-tVec(i-1))))/(2*a))*Z(i
           -1,:);
end

end
```

HullWhiteSwaption.m

```
function price = HullWhiteSwaption(t,T,X,rt,paymentTimes,f0,
   P0,a,sigma,optType)
%HULLWHITESWAPTION calculates the value of either a payer or
%receiver swaption with nominal of 1
%
%   t                   : value time
%   T                   : swaption maturity
%   X                   : strike
%   rt                  : current short rate
%   paymentTimes        : payment times of underlying swap
%   f0                  : initial inst. forward rate curve
%   P0                  : initial ZCB curve
%   optType             : 'receiver' or 'payer'

yrFracs = [paymentTimes(1)-T;diff(paymentTimes)];
c = [X*yrFracs(1:end-1); 1+X*yrFracs(end)];
obj = @(r) sum(c.*HullWhiteBondPrices(T,r,P0,f0,a,sigma,
   paymentTimes))-1;
rstar = fzero(obj,rt);
Xvec = HullWhiteBondPrices(T,rstar,P0,f0,a,sigma,
   paymentTimes);
Pt_T = HullWhiteBondPrices(t,rt,P0,f0,a,sigma,T);
Pt_S = HullWhiteBondPrices(t,rt,P0,f0,a,sigma,paymentTimes);

if strcmpi(optType,'payer');
    price = sum(c.*HullWhiteZBO(t,T,paymentTimes,Xvec,Pt_T,
```

```
        Pt_S,a,sigma,'put'));
elseif strcmpi(optType,'receiver');
    price = sum(c.*HullWhiteZBO(t,T,paymentTimes,Xvec,Pt_T,
        Pt_S,a,sigma,'call'));
end
```

HullWhiteZBO.m

```
function price = HullWhiteZBO(t,T,S,X,Pt_T,Pt_S,a,sigma,
   optType)
%HullWhiteZBO prices a put or call on a ZCB under the
%Hull-White model
%
%   t            : value date
%   T            : option maturity
%   S            : bond maturity (S>=T)
%   X            : strike
%   Pt_T, Pt_S : bond prices P(t,T) and P(t,S)
%   a            : mean-reversion rate
%   sigma        : short rate vol
%   optType      : 'put' or 'call'

B = @(t,T) (1-exp(-a*(T-t)))/a;
sigma_p = sigma*sqrt((1-exp(-2*a*(T-t)))/(2*a)).*B(T,S);
h = log(Pt_S./(X*Pt_T))./sigma_p+sigma_p/2;

if strcmpi(optType, 'call')
    price = Pt_S.*normcdf(h)-X.*Pt_T.*normcdf(h-sigma_p);
elseif strcmpi(optType,'put')
    price = X.*Pt_T.*normcdf(-h+sigma_p)-Pt_S.*normcdf(-h);
end

end
```

IRSpricer.m

```
function output = IRSpricer(t,zCurve,FSR,resetDates,
   paymentDates,fixings,out)
%IRSpricer returns either the fair-swap rate or theoretical
%price for a payer IRS with a notional of 1
%
%   t                 : value date
%   zCurve            : zcb curve at time t
%   FSR               : fair swap rate
%   resetDates        : swap reset times
%   paymentDates      : swap payment times
```

```
%   fixings          : floating cashflows which have become
%                      fixed
%   out              : required output

if strcmpi(out,'fsr');
    DF1 = interp1(zCurve(:,1),zCurve(:,2),resetDates-t);
    DF2 = interp1(zCurve(:,1),zCurve(:,2),paymentDates-t);
    yrFracs = paymentDates-resetDates;
    output = (DF1(1)-DF2(end))/(sum(yrFracs.*DF2));
    return
end

tmp = find(resetDates<=t,1,'last');
remResetDates = [resetDates(tmp); resetDates(resetDates>t)];
remPaymentDates = paymentDates(paymentDates>t);
remYrFracs = remPaymentDates - remResetDates;
DF1 = interp1(zCurve(:,1),zCurve(:,2),remResetDates-t);
DF2 = interp1(zCurve(:,1),zCurve(:,2),remPaymentDates-t);
vFix = sum(FSR*remYrFracs.*DF2);
vFloat = fixings(fixings(:,1)==remPaymentDates(1),2)*DF2(1)*
   remYrFracs(1)+...
    DF2(1)-DF2(end);
output = vFloat-vFix;
end
```

LGMbondPrices.m

```
function ZCB = LGMbondPrices(t,Xt,P0,alpha,H,Tvec)
%LGMBONDPRICES prices ZCBs in the LGM model at time t for
%maturities T1,T2,...TN
%
%   P0        : initial ZCB prices
%   alpha,H : params
%   t         : settlement
%   Tvec      : maturity dates

Pvec = interp1(P0(:,1),P0(:,2),Tvec);
Pt = interp1(P0(:,1),P0(:,2),t);
alpha2 = @(x) alpha(x).^2;
eta = integral(alpha2,0,t);

ZCB = (Pvec./Pt).*exp((H(t)-H(Tvec))*Xt+0.5*(H(t)^2-H(Tvec)
   .^2).*eta);

end
```

LGMrates.m

```
function paths = LGMrates(X0,alpha,tVec,N)
%LGMRATES simulates simulates N state process (Xt) paths
%under the LGM model at times t_1, t_2, ... t_n
%
%   X0      : initial value of the state process Xt
%   alpha   : instantaneous vol parameter
%   tVec    : vector of times t_0, t_1,...t_n
%   N       : number of paths to simulate

n = length(tVec)-1;
paths = [repmat(X0,1,N); zeros(n,N)];
Z = randn(n,N);

alpha2 = @(x) alpha(x).^2;
eta = arrayfun(@(t) integral(alpha2,0,t), tVec);
deta = diff(eta);

for i = 2:n+1
    paths(i,:) =  paths(i-1,:)+sqrt(deta(i-1))*Z(i-1,:);
end

end
```

LGMSwaption.m

```
function price= LGMSwaption(t,T,K,Xt,paymentTimes,P0,alpha,H
   ,optType)
%LGMSwaption calculates the price of a European swaption in
%the LGM model
%
%   t                 : value time
%   T                 : swaption maturity
%   K                 : strike
%   Xt                : current state variable
%   paymentTimes      : payment times of underlying swap
%   P0                : initial ZCB curve
%   alpha, H          : LGM params (functions of time)
%   optType           : 'receiver' or 'payer'

alpha2 = @(x) alpha(x).^2;

Pvec = LGMbondPrices(t,Xt,P0,alpha,H,[T;paymentTimes]);
Pt = LGMbondPrices(t,Xt,P0,alpha,H,t);
yrFracs = [paymentTimes(1)-T;diff(paymentTimes)];
```

```
Hvec = H(paymentTimes);
Ht = H(t);
HT = H(T);
zetaT = integral(alpha2,0,T);
zetat = integral(alpha2,0,t);
obj = @(y) sum(K*yrFracs.*Pvec(2:end).*exp(-(Hvec-HT)*y-0.5*
   zetaT*(Hvec.^2-...
    HT.^2)))+Pvec(end)*exp(-(Hvec(end)-HT)*y-0.5*zetaT*(Hvec
       (end).^2-HT.^2))-Pvec(1);
yStar = fzero(obj,0);

if strcmpi(optType,'payer');
    price = -sum(K*yrFracs.*Pvec(2:end).*(exp(-(Hvec-Ht)*Xt
       -0.5*(Hvec.^2-Ht.^2)*zetat)/Pt).*(normcdf((yStar-Xt+
       Hvec*(zetaT-zetat))/sqrt(zetaT-zetat))))+...
        -Pvec(end)*(exp(-(Hvec(end)-Ht)*Xt-0.5*(Hvec(end)
           .^2-Ht.^2)*zetat)/Pt)*(normcdf((yStar-Xt+Hvec(end
           )*(zetaT-zetat))/sqrt(zetaT-zetat)))...
        +Pvec(1)*(exp(-(HT-Ht)*Xt-0.5*(HT.^2-Ht.^2)*zetat)/
           Pt)*(normcdf((yStar-Xt+HT*(zetaT-zetat))/sqrt(
           zetaT-zetat)));
elseif strcmpi(optType,'receiver');
    price = sum(K*yrFracs.*Pvec(2:end).*(exp(-(Hvec-Ht)*Xt
       -0.5*(Hvec.^2-Ht.^2)*zetat)/Pt).*(1-normcdf((yStar-Xt
       +Hvec*(zetaT-zetat))/sqrt(zetaT-zetat))))+...
        Pvec(end)*(exp(-(Hvec(end)-Ht)*Xt-0.5*(Hvec(end).^2-
           Ht.^2)*zetat)/Pt)*(1-normcdf((yStar-Xt+Hvec(end)
           *(zetaT-zetat))/sqrt(zetaT-zetat)))...
        -Pvec(1)*(exp(-(HT-Ht)*Xt-0.5*(HT.^2-Ht.^2)*zetat)/
           Pt)*(1-normcdf((yStar-Xt+HT*(zetaT-zetat))/sqrt(
           zetaT-zetat)));
end

end
```

LSMpaths.m

```
function [paths, regCoeffOut] = LSMpaths(tVec,rFactorPaths,
   numerairePaths,VT,cfTimes,cf,basis,...
    regCoeff)
%LSMPATHS estimates the price paths of a contract by
%regressing 'raw continuation' values one step ahead against
%realised values of the risk factor.
%
%   rFactorPaths    : simulated risk factor paths
%   numerairePaths : numeraire paths associated with
```

```
%                          simulated risk factor paths
%   VT                   : terminal value of the derivative on
%                          each risk factor path
%   basis                : basis functions specified in a cell
%                          array
%   regCoeff             : out-of-sample regression coefficients

[n, N] = size(rFactorPaths);
paths = [zeros(n-1,N);VT];
regCoeffOut = zeros(size(basis(1),2),n-1);

for i = n-1:-1:1
    if any(tVec(i+1) == cfTimes)
        Y = ((paths(i+1,:)+cf(tVec(i+1)==cfTimes,:)).*
           numerairePaths(i,:)./numerairePaths(i+1,:)).';
    else
        Y = (paths(i+1,:).*numerairePaths(i,:)./
           numerairePaths(i+1,:)).';
    end
    X = cell2mat(basis(rFactorPaths(i,:).'));
    if isempty(regCoeff)
        beta = ((X.'*X)^-1)*X.'*Y;
        paths(i,:) = X*beta;
        regCoeffOut(:,i) = beta;
    else
        paths(i,:) = X*regCoeff(:,i);
    end
    plot(rFactorPaths(i,:).',Y,'.k',rFactorPaths(i,:).',
       paths(i,:),'.r')
end

end
```