

A stochastic monetary policy interest rate model

Claudio Albanese
claudio@level3finance.com

Manlio Trovato
Merrill Lynch
manlio_trovato@ml.com

May 3, 2007

Abstract

We present a three factor interest rate term structure model solved on a continuous-time lattice and constructed with local volatility, asymmetric jumps, stochastic volatility regimes and stochastic monetary policy.

We take the 3-month spot LIBOR rate as modelling primitive. An analysis of the historical data suggests that the process is subject to drift regimes. As the 3-month spot LIBOR rate is not an asset price process, the drift is not constrained by the no-arbitrage condition. We incorporate a direct specification of the drift process in our model: this is a novel approach in interest rate modelling and it represents one of the main contributions of this work. The model can be formally written as $dL_t = \mu_{a_t}(L_t)dt + \sigma_{b_t}(L_t)dW_t + \text{jumps}$ where the drift and the volatility terms are stochastic and driven by the processes a_t, b_t , correlated to the rates themselves. We show that the model can achieve a persistent smile structure across maturities with a nearly time homogeneous parameterisation, in good qualitative agreement with the EUR market. We also show that the model is able to explain most yield curve shapes of interest, and that these are crucially affected by the drift regimes, as one would intuitively expect. Furthermore we explain that drift modelling provides a powerful tool in order to control the long term behaviour of the process without affecting considerably the dynamics at earlier times. In fact, whilst jumps are predominant at short maturities, stochastic volatility has the greatest impact at medium maturities, whereas drift is the main driver in the long end. In addition, we argue that direct drift modelling is economically meaningful and allows one to impose economic views on central banks monetary policy into the model.

1 Introduction

In this work we present a novel approach to interest rate modelling, which incorporates stochastic monetary policy. Taking inspiration on short rate models, we choose a modelling primitive which is not an asset price process. The main consequence of this choice is that the drift term can be freely specified as this

is not implied by the specification of the volatility term and the no-arbitrage condition. We exploit this property and enhance modelling flexibility by introducing an additional degree of freedom and allowing direct control on the drift process.

With the notable exception of short rate models, this would not be possible in most models in the traditional literature: in this case an asset price process, or a deterministic function of it, is taken as a modelling primitive and financial modelling is therefore limited to the specification of the volatility term.

Whilst short rate modelling provides a suitable mathematical framework for expressing economic views, like for example the mean reversion effect in rates, the need for analytical tractability has severely limited its application domain. We attempt to overcome its limitation and propose an interest rate model built upon the specification of a conditional local volatility process, with asymmetric jumps introduced by subordination, a volatility regime process and a drift regime process.

We enforce a nearly time-homogeneous parameterisation and show that the model is able to generate a persistent smile across tenors and maturities, in good qualitative agreement with the EUR market. We also show that the model is able to explain most yield curve shapes of interest and find that these shapes are considerably affected by the volatility regime process, as one would intuitively expect.

We take the 3-month LIBOR spot rate L_t as our modelling primitive. The term structure of discount bonds is recovered through a model dependent formula:

$$Z(t_i, t_j) = \mathbb{E}_{t_i} \left[\prod_{k=i}^{j-1} \frac{1}{1 + \tau L_{t_k}} \right] \quad (1)$$

where the expectation is taken in the spot rolling measure with simple compounding, L_{t_k} denotes the prevailing 3-month spot LIBOR rate resetting at time t_k and τ is the quarterly accrual fraction ($\tau = 0.25$). As L_t is not a fixed asset price process, the no-arbitrage condition does not impose any restrictions on its drift. The process for L_t can formally be written as follows:

$$dL_t = \mu_{a_t}(L_t)dt + \sigma_{b_t}(L_t)dW_t + \text{jumps} \quad (2)$$

where W_t is a one-dimensional Brownian motion and the drift and the volatility terms are stochastic and driven by the processes a_t, b_t correlated to the rates themselves. Both the drift and the volatility terms can be freely specified, hence the specification of the process for the drift represents an additional degree of freedom available to the modeller. We show how the drift regimes can be used as a direct control for the calibration of the long dated portion of the volatility smile in high and low strike regions without considerably affecting the nature of the process at earlier maturities. In fact, it is well known that jumps are predominant at short maturities and stochastic volatility persists for longer time

span. Our finding is that the drift process, associated with the first moment of the L_t process, has a predominant impact at long maturities. The three different components can therefore be effectively used to drive the process at different time spans whilst retaining a time homogeneous parameterisation.

Our model is solved by means of discretisation schemes based on continuous-time Markov chains. Discretisation schemes of similar type have been previously discussed in (Albanese and Kuznetsov, 2005) and (Albanese and Kuznetsov, 2003), and applied to models for which the spectrum of the Markov generator can be computed in analytically closed form. While retaining the insights provided by the spectral analysis treatment in the previous papers, here we renounce to analytic solvability and follow instead a non-parametric approach, similarly to what we have proposed in previous works as (Albanese and Trovato, 2005). The new idea is to leverage not on the ability to evaluate special functions but instead on numerical linear algebra routines in order to compute a matrix exponentiation. Our discretization schemes are based on operator methods and level-3 numerical linear algebra (i.e. matrix-matrix multiplication), which is considerably more efficient than traditional finite difference schemes, limited to level-2 BLAS (i.e. matrix-vector multiplication).

The introduction of stochastic monetary policy in interest rate modelling is the main technical contribution of this work. In this framework, the modeller has the flexibility of controlling the drift process. The impact on derivative pricing can be assessed and the comparison with other established modelling framework may help the modeller gaining more insight on the factors driving pricing functions and risk patterns. Furthermore, the use of operator methods and continuous-time lattices also represents a distinctive feature of this modelling framework.

The sequel is organised as follows: we first discuss the construction of continuous-time lattices and their convergence properties for the particular case of a conditional local volatility process; then we describe the specification of the interest rate model and its calibration. Finally, we present an application to callable CMS spreads and show the impact of drift regimes on the price functions.

2 Continuous-time lattices

2.1 Construction of continuous-time lattices

The specification and solution of our model is defined by constructing a continuous-time Markov chain on a lattice with finitely many states. In this section we discuss the basic concepts underlying the construction of the lattice. Although our model is not limited to a simple diffusion process, it is convenient to present the methodology of the construction of a continuous-time lattice for this simple case. We take as a modelling primitive a generic asset price process F_t and assume that it follows a diffusion process:

$$dF_t = \mu(F_t)dt + \sigma(F_t)dW_t \quad (3)$$

where the drift $\mu(F_t)$ and the local volatility $\sigma(F_t)$ are bounded and continuous functions defined on the domain of F_t . It is well known (see for example (Karatzas and Shreve, 1991)) that the Markov generator \mathcal{L}_E , defined on a continuum state space E , for a diffusion process F_t acts on a twice differentiable function u as follows:

$$(\mathcal{L}_E u)(F_t) = \mu(F_t) \frac{\partial u}{\partial F}(F_t) + \frac{1}{2} \sigma^2(F_t) \frac{\partial^2 u}{\partial F^2}(F_t). \quad (4)$$

We consider the probability density function $u(F_0, t; dF, T)$, which is known to satisfy the backward Kolmogorov equation:

$$\frac{\partial u}{\partial t} + \mathcal{L}_E u = 0 \quad (5)$$

with boundary condition $u(F_0, t; dF, T) = \delta(F_T - dF)$. The solution to eq. (5) can be formally written as:

$$u(F_0, t; dF, T) = e^{\mathcal{L}_E(T-t)}(F_0, dF). \quad (6)$$

It is therefore obvious that, in order to solve for the probability density function u , also referred to as pricing kernel, all we need is the specification of the Markov generator \mathcal{L}_E . We therefore proceed with the construction of the approximating Markov chain that discretises the local volatility diffusion process. The construction is based on the specification of the local properties of the Markov chain, i.e. on the specification of the Markov generator on finitely many states. It can be proved that, if the Markov generator satisfies certain local consistency conditions, then the resulting Markov chain is a correct approximation of the original process (see for example (Kushner and Dupuis, 2001)), meaning that the discretised process converges to the continuous limit process in a weak or distributional sense. Once the Markov generator is correctly specified, the transition probabilities between any two points in the continuous-time lattice can be computed by calculating the exponential of the Markov generator matrix itself.

We define a discrete domain for the forward rate process F_t and build a continuous-time Markov chain with finitely many states. Let our state space be $\Omega = \{0, 1, 2, \dots, N\}$, the set of the first N integers together with zero. Let $F : \Omega \rightarrow \mathbb{S}$ be a map from Ω to some countable set $\mathbb{S} \in \mathbb{R}$, such that $F(x) \leq F(y)$ for all $x \in \Omega$ and $y \in \Omega - \{0\}$, such that $x < y$. The map $F : \Omega \rightarrow \mathbb{S}$ defines the discrete grid for all possible values that the discretized version $F_{\Omega t}$ of the process F_t can take on the lattice Ω . We note that the grid is not necessarily uniform, therefore it is possible to adopt a judicious and efficient node placing scheme. In particular, for a given total number of nodes N , one can improve the convergence of the discretized process by reducing the node spacing in region closer to a guessed expected value of F , or alternatively around the spot value F_0 , and making use of a sparser grid in regions of very high and low levels in F .

It can be shown (see (Albanese and Mijatovic, 2006) and (Kushner and Dupuis, 2001)) that an appropriate Markov chain approximation of the diffusion process can be constructed by imposing the following local consistency conditions:

$$\mathbb{E}_t [F_{t+dt} - F_t] = \mu dt \quad (7)$$

$$\mathbb{E}_t \left[(F_{t+dt} - F_t)^2 - \mathbb{E}_t [F_{t+dt} - F_t]^2 \right] = \sigma^2 dt \quad (8)$$

Let \mathcal{L} be a finite difference discretisation of the operator \mathcal{L}_E (see (Albanese and Mijatovic, 2006) and (Kushner and Dupuis, 2001)):

$$\mathcal{L} := \mu \nabla_h + \frac{\sigma^2}{2} \Delta_h \quad (9)$$

where ∇_h and Δ_h are the discrete finite difference and Laplace operator respectively, defined for a lattice spacing $h = F(x) - F(x-1)$, with $x \in \Omega - \{0\}$. Then one can write the local consistency conditions in terms of the discretised Markov generator \mathcal{L} :

$$\sum_y \mathcal{L}(x, y)(F(y) - F(x)) = \mu(F(x)) \quad (10)$$

$$\sum_y \mathcal{L}(x, y)(F(y) - F(x))^2 = \sigma^2(F(x)). \quad (11)$$

The above two equations, together with the probability conservation condition:

$$\sum_y \mathcal{L}(x, y) = 0 \quad (12)$$

uniquely identify the Markov generator for the diffusion process.

Once the Markov generator for the approximating Markov chain is known, the transition probability kernel defined on the lattice can be obtained by computing the exponential of the Markov generator matrix itself:

$$u(x, t; y, T) = e^{\mathcal{L}(T-t)}(x, y). \quad (13)$$

2.2 Bochner subordinators

Bochner subordinators are a special type of time changed processes and were first introduced by (Bochner, 1955). By means of Bochner subordinators, one can efficiently introduce jumps, independent from the subordinated process, in the model construction. In this section, for convenience, we work first with stochastic process of continuous random variables and then define the discrete space approximation.

Definition 1 *The process T_t is called a subordinator if it is a right-continuous non decreasing process with values in \mathbb{R}^+ such that $T_0 = 0$ and it has independent and homogeneous increments.*

A subordinator is associated to a measure $\nu_t(x)$ such that, for all continuous functions of compact support $f(x)$, the following holds:

$$\mathbb{E} \left[\int_0^\infty f(T_s) ds \right] = \int_0^\infty f(x) \nu_t(dx). \quad (14)$$

The following theorem characterises Bochner subordinators and it provides the basis for an efficient construction of a Markov generator of a subordinated conditional local volatility process.

Theorem 2 *A subordinator is a Bochner subordinator if and only if:*

$$\int_0^\infty e^{-\lambda x} \nu_t(dx) = e^{-t\phi(\lambda)}. \quad (15)$$

for some function $\phi(\lambda)$ called the *Bernstein function*.

It is possible to show that (see (Albanese and Kuznetsov, 2003)), if \mathcal{L}_E is the Markov generator of a diffusion process, defined on a continuum state space E , and $\phi(\lambda)$ is the *Bernstein* function of a subordinator T_t , then the generator of the subordinated process \mathcal{L}_E^j is simply given by:

$$\mathcal{L}_E^j = -\phi(-\mathcal{L}_E). \quad (16)$$

Let \mathcal{L} be the discretisation of the Markov generator \mathcal{L}_E and defined on a finite lattice, then the following algorithm may be used in order to apply the Bochner subordinator to the discretised process. First, the matrix \mathcal{L} is diagonalised by means of numerical routines like *geev* of LAPACK:

$$\mathcal{L} = U\Lambda U^{-1} \quad (17)$$

where Λ is a diagonal matrix containing the eigenvalues (i.e. the spectrum) of \mathcal{L} on its main diagonal and U is the matrix having as columns the right eigenvectors; then one can apply the *Bernstein* function to the matrix \mathcal{L} by taking advantage of the following fundamental result of functional calculus:

Theorem 3 *If $\mathcal{L} \in \mathbb{C}^{n \times n}$, with $\mathcal{L} = U\Lambda U^{-1}$ and $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$, then a function φ , defined on the spectrum of \mathcal{L} , can be applied to the matrix \mathcal{L} by means of the following formula:*

$$\varphi(\mathcal{L}) = U\varphi(\Lambda)U^{-1} \quad (18)$$

As Λ is a diagonal matrix, the calculation of $\varphi(\Lambda)$ is a very simple task:

$$\varphi(\Lambda) = \begin{pmatrix} \varphi(\lambda_1) & 0 & \cdots & 0 & 0 \\ 0 & \varphi(\lambda_2) & \cdots & 0 & 0 \\ 0 & 0 & \varphi(\lambda_3) & 0 & 0 \\ \vdots & \vdots & 0 & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & \varphi(\lambda_n) \end{pmatrix} \quad (19)$$

It should be noted that the theory of subordination is defined on a continuous state space. One may therefore want to show that the subordinated discretised process admits a continuous limit and that the resulting discretised process converges to this limit in distributional sense. We don't pursue this further, but we note that a similar proof has been provided for an alternative construction of the approximating Markov chain for a jump diffusion model in (Kushner and Dupuis, 2001).

2.3 Convergence properties

Two types of convergence properties have been shown to hold for the approximating Markov chain constructed in the previous section: convergence of the finite dimensional distribution and weak convergence, or convergence in distribution.

For the former case, it has been shown that the transition probability kernel of the discretised process converges to the probability density function of the diffusion process, when the discretisation step h tends to zero (see (Albanese and Mijatovic, 2006)). The convergence rate has been shown to be $O(h^2)$.

For a slight different Markov chain approximation, where both the state space and time are discretised, (Kushner and Dupuis, 2001) show that the discretised Markov chain converges *in distribution* to the diffusion process. This is a stronger result, as it implies the convergence between probability measures, however in this case convergence rate estimations are not obtainable.

In both cases, it has been shown that, if a central difference approximation is used for the discrete first difference operator, then the following condition must hold:

$$\inf_x [\sigma^2(F(x)) - h |\mu(F(x))|] \geq 0 \quad (20)$$

where h is the lattice spacing and $x \in \Omega$.

2.4 Matrix exponentiation

The computation of the discretised transition probability kernel boils down to the valuation of the exponential of the Markov generator matrix. In the discretisation schemes discussed in (Albanese and Kuznetsov, 2003), this was performed uniquely by an analytical spectral decomposition of the Markov generator. This

represented a limitation, as the analytical spectral decomposition is possible only for a limited range of processes. Instead we follow a non parametric approach and propose to perform the exponentiation numerically. This simple observation has an important impact for the modeller, as it allows for a considerably enhanced modelling flexibility. Obviously, the drawback being the fact that the method relies on fast and accurate numerical methods.

A comprehensive review on available methods for the computation of the exponential of a matrix is (Moler and Loan, 2003). Unfortunately there is no single method that appears to work well in all situations and the classification criterion for selecting the most appropriate method is not straightforward. In our experience we have seen that, generally speaking, for reasonably symmetric matrices the diagonalisation method works well, even in double precision arithmetic. However, for asymmetric matrices this method becomes ill-conditioned and suffers from the pseudo-spectra problem. In this case, the scaling and squaring method appears to be more reliable. For application of our interests, the calculation of the transition probability kernel for simple models based on diffusion processes can be safely solved with the diagonalisation technique, and in general the scaling and squaring method is more stable and therefore represents our preferred choice. We now review these two methods in the following sections.

2.4.1 Matrix diagonalisation

Given a discretized Markov generator \mathcal{L} defined on a finite lattice Ω containing N points, we consider the following pair of eigenvalue problems:

$$\mathcal{L}u_n = \lambda_n u_n \qquad \mathcal{L}^T v_n = \lambda_n v_n \qquad (21)$$

where the superscript T denotes matrix transposition, u_n and v_n are the right and left eigenvectors of \mathcal{L} , respectively, whereas λ_n are the corresponding eigenvalues. Except for the simplest cases, the Markov generator \mathcal{L} is not a symmetric matrix, hence u_n and v_n are different. Also, in general, the eigenvalues are not real. We are only guaranteed that their real part is non-positive $\text{Re}\lambda_n \leq 0$ and that complex eigenvalues occur in complex conjugate pairs, in the sense that if λ_n is an eigenvalue then its conjugate $\bar{\lambda}_n$ is also an eigenvalue.

The diagonalization problem can be rewritten in the following matrix form:

$$\mathcal{L} = U\Lambda U^{-1} \qquad (22)$$

where U is the matrix having as columns the right eigenvectors and Λ is the diagonal matrix having the eigenvalues λ_i as elements. The numerical implementation of the diagonalization problem can be efficiently carried out by means of routines such as `geev` in LAPACK, however care must be taken in analysing the stability and accuracy of the spectrum. The method turns out to be unstable, in double precision, when the conditioning number $\text{cond}(V) = \|U\| \|U^{-1}\|$ of the matrix of eigenvector is small. In our experience this is more likely to happen when the Markov generator is asymmetric. It is possible to evaluate

the circles of uncertainty for each eigenvalue (also called pseudo-spectrum) by means of perturbation techniques on the Markov generator or with alternative algorithms, as described in (Threfethen, 1999).

Once the matrix has been decomposed in eigenvalues and eigenvectors, the exponential of the matrix can be computed as follows:

$$e^{t\mathcal{L}} = Ue^{t\Lambda}U^{-1} \quad (23)$$

where $e^{t\Lambda}$ is trivially evaluated:

$$e^{t\Lambda} = \begin{pmatrix} e^{t\lambda_1} & 0 & \dots & 0 & 0 \\ 0 & e^{t\lambda_2} & \dots & 0 & 0 \\ 0 & 0 & e^{t\lambda_3} & 0 & 0 \\ \vdots & \vdots & 0 & \ddots & \vdots \\ 0 & 0 & \dots & 0 & e^{t\lambda_n} \end{pmatrix} \quad (24)$$

2.4.2 Fast exponentiation

The scaling and squaring method, also known as *fast exponentiation method*, is based on the following fundamental property of the exponential function:

$$e^{t\mathcal{L}} = (e^{\delta t\mathcal{L}})^{\frac{t}{\delta t}} \quad (25)$$

Let us choose $\delta t > 0$ as the largest time interval for which both of the following properties are satisfied:

$$(FE1) \quad \min_{y \in \Lambda} (1 + \delta t\mathcal{L}(y, y)) \geq 1/2$$

$$(FE2) \quad \log_2 \frac{t}{\delta t} = n \in \mathbb{N}.$$

To compute $e^{t\mathcal{L}}(x, y)$, we first define the elementary propagator

$$u_{\delta t}(x, y) = \delta_{xy} + \delta t\mathcal{L}(y, y) \quad (26)$$

and then evaluate in sequence $u_{2\delta t} = u_{\delta t} \cdot u_{\delta t}$, $u_{4\delta t} = u_{2\delta t} \cdot u_{2\delta t}$, ... $u_{2^n \delta t} = u_{2^{n-1} \delta t} \cdot u_{2^{n-1} \delta t}$.

The repeated matrix-matrix multiplication can be efficiently calculated with Level-3 BLAS numerical routines. Due to processor-specific optimisation, clever cache management and hardware acceleration on GPUs, these operations can nowadays be performed at great speed on low cost hardware.

3 The interest rate model

3.1 Desiderata for an interest rate term structure smile model

The large variety of payoffs available in the interest rate derivatives market means that trading desks build up exposure not only to interest rate movements,

but also to a few other derived state variables, some of which are not easily observable on the market. For consistent pricing and robust hedging one would want to try and make use of the same model across all products, and retain efficient numerical solubility and stable hedge ratios computation. Whilst we do not engage in a discussion on what should be the selection criterion for a good interest rate model, which has been an active debate in the literature, we note that an interest rate derivatives desk of a major financial institution is likely to build up exposure to:

- interest rates movements;
- at-the-money implied volatility movements;
- movements of the liquid portion of the smile (i.e. the relationship between the implied volatility and strikes);
- movements of the asymptotic smile, i.e. in regions of extreme strike levels;
- forward volatility movements;
- forward rates correlation movements.

Interest rates and at-the-money implied volatilities are observable in the market, as the liquid portion of the smile. The latter is driven by European swaption contracts and, as an example, for the current EUR market it spans a strike range between 2% and 8% approximately. The asymptotic smile spans extreme strike levels and is driven primarily by constant maturity swaps (CMS). These exotic structures receive a floater, the coupon being a spread over LIBOR, and pay the equilibrium swap rate of a fixed tenor prevailing at each coupon date. An analysis of the CMS leg coupon structure leads to the conclusion that CMS contracts are particularly sensitive to the asymptotic behavior of implied volatilities for very large strikes. Therefore market CMS rates are an average indication of the asymptotic smiles and they suggest that implied volatilities flatten out and converge asymptotically to a constant. Forward volatility is not directly observable in the interest rates market, a part from some fairly illiquid products, and implied forward rates correlation is only observable for a few forward rate pairs over a limited range of strikes.

A consistent model aims to recover all market prices and observables and can be used to extrapolate the values for those quantities that, although not directly observable on the market, play an important role in the pricing and risk management of some interest rate derivative products. In the construction of our model we make every effort to retain the economic interpretation of the assumptions imposed on the process. The model is built up from a simple conditional diffusion process and then enhanced with the addition of other more complicated effects, including jumps, volatility and drift regimes.

The first goal of our model construction is to try and describe consistently both the liquid and the asymptotic portion of the smile, in good qualitative

agreement with the market. The rich model parameterization and the ability to control the transition probability of the spot rate to span and stay in different value regions allow us to achieve this goal.

The second goal is to formulate a robust model with a parameter set that does not need to change often in time. For this we use time-homogeneity as a criterion guiding our methodology. Time-homogeneity also help achieving the third goal, which is to ensure that the model implies realistic values for the implied forward volatility. In fact, in a time-homogeneous model the implied forward smile is necessarily similar to the spot smile. This is a desirable feature for an interest rate model.

The fourth goal is to allow to recover the implied or historical correlation structure between forward rates. The drift regime process provides the necessary control for this.

3.2 The modelling framework

We define a tenor structure t_i , $i \in \{0, 1, \dots, M\}$ such that $\tau = t_j - t_{j-1} = 0.25$, for each $j \in \{1, \dots, M\}$. We then take the 3-month LIBOR spot rate L_t as our modelling primitive. The term structure of discount bonds is recovered through a model dependent formula:

$$Z(t_i, t_j) = \mathbb{E}_{t_i} \left[\prod_{k=i}^{j-1} \frac{1}{1 + \tau L_{t_k}} \right] \quad (27)$$

where the expectation is taken in the risk neutral measure, $Z(t_i, t_j)$ is the value at t_i of a zero coupon bond paying one currency unit at t_j and L_{t_k} is the prevailing 3-month LIBOR spot rate at time t_k . As L_t is not a fixed asset price process, the no-arbitrage condition does not impose any restrictions on its drift. The process for L_t can formally be written as follows:

$$dL_t = \mu_{a_t}(L_t)dt + \sigma_{b_t}(L_t)dW_t + \text{jumps} \quad (28)$$

where the drift and the volatility terms are stochastic and are driven by the processes a_t, b_t correlated to the rates themselves. Both the drift and the volatility terms can be freely specified for the construction of the process for the 3-month LIBOR rate.

In our modelling framework the process for the LIBOR rate is driven by a conditional local volatility process with asymmetric jumps, a volatility regime variable process and a drift regime variable process. The conditional local volatility is specified as a CEV process, where the exponent is made a function of the rates themselves. Asymmetric jumps are also introduced, in order to describe the pronounced and asymmetric smile at very short maturities. The volatility regime variable is constrained to take three possible values: **low**, **medium** and **high** volatility. The drift regime variable can take values corresponding to regimes of **rising**, **stable**, **falling** rates and **deflation**.

In this modelling framework, the yield curve dynamics is driven by three factors: the conditional process for the 3-month LIBOR spot rate, the process for the volatility regime variable and the process for the drift regime variable.

3.3 Historical time series analysis

We analyse the historical time series for 3-month LIBOR spot rate for three major currencies: EUR, USD and JPY (see fig.(1)). It is obvious that the 3-month LIBOR spot rate is driven by trends which may last for periods that span multiple years. One can easily identify falling, stable and rising trends, in the EUR and USD time series. In our model, these trends are associated with drift regimes corresponding to falling, stable and rising rates. The JPY time series is interesting because it spans regions not explored by the EUR and USD time series, and in particular regions where the rate is close to zero. When the rate approaches these extremely low levels, it is subject to deflationary pressure and it may be constrained to stay at such low levels for an extended period of time, as it has happened for the Japanese economy between 2000 and 2005. In our model, we include a deflationary drift regime which is associated to this phenomenon.

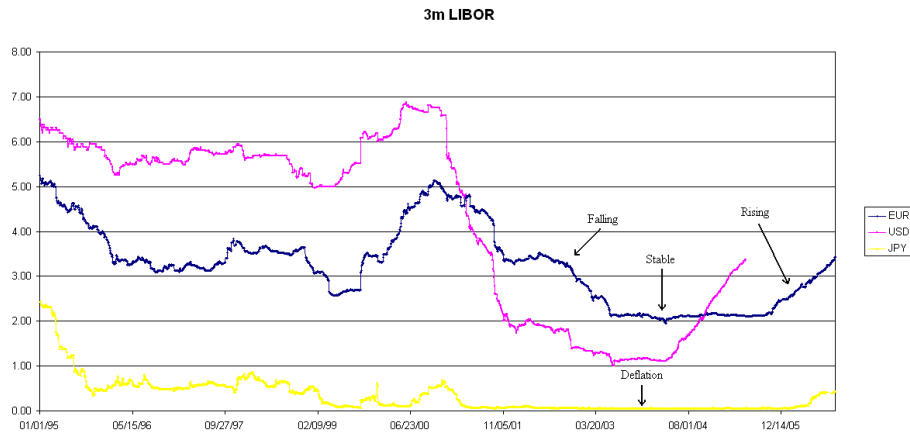


Figure 1: spot 3m Libor rate historical time series

3.4 The model construction

3.4.1 The conditional local volatility processes

We introduce $N_L = 70$ states for the conditional local volatility process, $N_a = 3$ states for the volatility regime variable and $N_b = 4$ states for the drift regime

variable. Let Ω_L be a finite set $\{0, 1, \dots, N_L - 1\}$ containing the first $N_L - 1$ integers together with zero, Ω_a the finite set $\{0, 1, 2\}$ and Ω_b the finite set $\{0, 1, 2, 3\}$. Let $L : \Omega_L \rightarrow \mathbb{S}$ be an increasing positive function such that $L(0) = 0$, where \mathbb{S} is some countable set of \mathbb{R} . Also, let $y_i = (x_i, a_i, b_i)$ be a generic element of the lattice $\Omega = \Omega_L \times \Omega_a \times \Omega_b$, where $x_i \in \Omega_L, a_i \in \Omega_a$, and $b_i \in \Omega_b$.

Conditional to a volatility regime state a and a drift regime state b , we define an approximating Markov chain associated with the conditional local volatility process for the 3-month LIBOR spot rate. The Markov generator is constructed by imposing probability conservation and the local consistency properties, in the sense of eq. (12), (10) and (11) respectively. For the conditional local volatility process, the drift is constructed with a series of piecewise linear functions of the underlying rate, as in fig. (2(a)), and the local volatility is of the form:

$$\sigma_b(L(x)) = \text{Max}(\bar{\sigma}_b L(x)^{\beta(L(x))}, \sigma_{\max}). \quad (29)$$

The local volatility function is of a *CEV* type and it is constrained to be less than a maximum level in order to avoid probability accumulation at the boundary of the domain. The exponent is made dependent of the level of the underlying rates as follows:

$$\beta(L(x)) = \frac{A + L(x)}{B + L(x)} \quad (30)$$

such that $\beta(0) = \beta_0$ and $\beta(L_{ref}) = \beta_{ref}$. Fig. (3) and fig. (2(b)) report the beta function and the lognormal local volatility functions respectively. In the classical SDE formalism, the local volatility process, conditional on a drift regime a and a volatility regime b , can be written as follows:

$$dL_t = \mu_a(L_t)dt + \sigma_b(L_t)dW_t \quad (31)$$




Figure 2: Specification of drift and local volatility, conditioned on the drift and volatility regimes respectively




Figure 3: Specification of beta

3.4.2 Introducing jumps

At this stage of the construction one has the option to add jumps. Although in the examples discussed in this paper we are mostly focused on long dated callable exotics for which we find that the impact of jumps can be safely ignored, adding jumps involves negligible additional complexities and is thus worth considering and implementing in other situations. To add jumps, uncorrelated to the conditional local volatility process, we use the Bochner subordinators with the Bernstein function associated to the variance gamma process which received much attention in finance:

$$\phi(\lambda) = \frac{\mu^2}{\nu} \log \left(1 + \lambda \frac{\nu}{\mu} \right) \quad (32)$$

where μ is the mean rate and ν is the variance rate of the variance gamma process. We then follow the following procedure which accounts for the need to assign different intensities to up-jumps and down-jumps. The generator of the jump process can be expressed as the operator $-\phi(-\mathcal{L}_{\Omega_L})$. To produce asymmetric jumps, we specify the two parameters in (32) differently for the up and down jumps and compute separately two Markov generators

$$\mathcal{L}_{\Omega_L}^j_{\pm} = -\phi_{\pm}(-\mathcal{L}_{\Omega_L}) = -U_{\pm} \phi_{\pm}(-\Lambda) U_{\pm}^{-1}$$

where:

$$\phi_{\pm}(\lambda) = \frac{\mu_{\pm}^2}{\nu_{\pm}} \log \left(1 + \lambda \frac{\nu_{\pm}}{\mu_{\pm}} \right) \quad (33)$$

The new generator for our process with asymmetric jumps is obtained by combining the two generators above:

$$\mathcal{L}_{\Omega_L}^j = \begin{pmatrix} 0 & \cdots & \cdots & \cdots & 0 \\ \mathcal{L}_{\Omega_L-}^j(2,1) & d(2,2) & \mathcal{L}_{\Omega_L+}^j(2,3) & \cdots & \mathcal{L}_{\Omega_L+}^j(2,n) \\ \vdots & \vdots & \ddots & \cdots & \vdots \\ \mathcal{L}_{\Omega_L-}^j(n-1,1) & \mathcal{L}_{\Omega_L-}^j(n-1,2) & \cdots & d(n-1,n-1) & \mathcal{L}_{\Omega_L+}^j(n-1,n) \\ 0 & 0 & \cdots & \cdots & 0 \end{pmatrix}. \quad (34)$$

Here the element of the diagonal are chosen in such a way to satisfy probability conservation:

$$d(x_1, x_1) = - \sum_{x_2 \neq x_1} \mathcal{L}_{\Omega_L}^j(x_1, x_2) \quad (35)$$

Also notice that we have zeroed out the elements in the matrix at the upper and lower boundary: this ensures that there is no probability leakage in the process.

At this stage of the construction, we have therefore obtained a generator $\mathcal{L}_{\Omega_L}^j$ for the 3-month spot LIBOR rate process, whose dynamics is characterized by a combination of state dependent local volatility and asymmetric jumps. We note that the addition of jumps has not increased the dimensionality of the problem and is therefore computationally efficient.

3.4.3 Modelling the dynamics of stochastic volatility regimes

Stochastic volatility is modelled by means of volatility regimes. The volatility is allowed to switch between three possible states: **low**, **medium** and **high**. Evidence that the volatility process experiences patterns which can be effectively described by means of regimes has been previously discussed by (Rebonato and Joshi, 2001).

The transition probabilities between regimes, in our parameterisation choice, are set to be independent of the rate. The Markov generator has been constructed both by ensuring that transition probabilities between regimes reflect a meaningful view on the volatility process and by performing a fit to the volatility cube. The values used in our model parameterisation are the following:

$$\mathcal{L}_{\Omega_a} = \begin{pmatrix} -2.5 & 2.5 & 0 \\ 0.2 & -0.6 & 0.4 \\ 0 & 0.4 & -0.4 \end{pmatrix} \begin{array}{l} \leftarrow \textit{high} \\ \leftarrow \textit{medium} \\ \leftarrow \textit{low} \end{array}$$

The process on the regimes is typically much slower than the process on the state variable L_t , and it is responsible for driving the medium to long term behaviour of the term structure of the yield curve. The rate is allowed to transition from one regime to another with low frequency and we have chosen to allow only transition between contiguous regimes. Also, the process reverts from the **high** to the **medium** volatility state faster than any other regime transition.

3.4.4 Modelling the dynamics of a stochastic monetary policy

The elements of the Markov generator of the drift regime process are shown in fig. (4). These are defined to be a function of the underlying LIBOR rate. This functional dependence introduces a correlation between rates and drift regimes and allows to impose qualitative economic views. A few observations drive the specification of the elements of the Markov generator for the drift process:

- the drift regime process is allowed to jump with non-zero probability only between contiguous regimes;
- a part from the deflationary regime, beyond a certain rate level the rate may switch between different regimes with equal probability; therefore the transition probability intensities from one regime to another have the same value;
- falling to deflation: at very low rate levels, there is a large probability of switching from a falling to a deflation regime; whereas for increasing rate levels this probability rapidly decreases to zero;
- deflation to falling: at very low level there is a small probability of switching from a deflation to a falling regime; this probability rapidly increases with increasing level of rates;
- stable to falling: the probability of switching from a stable to a falling regime increases up to a constant as rates increase up to a reference level;
- falling to stable: the probability of switching from a falling to a stable regime decreases up to a constant as rates increase up to a reference level;
- stable to rising: the probability of switching from a stable to a rising regime decreases up to a constant as rates increase up to a reference level;
- rising to stable: the probability of switching from a rising to a stable regime increases up to a constant as rates increase up to a reference level.

One should note that the above behaviour is qualitatively consistent with the mean reversion phenomenon observed in the interest rates market. The definition of the Markov generator for the drift process is an important tool for the implementation of one's economic view into the model. Therefore, many possible variations may be specified, which may work just as well. What is important is that model provides a natural mathematical framework for such specification. The specification of such framework is one of the most important contributions of this work.

3.4.5 The combined process for the spot LIBOR rate

We complete the construction of our interest rate model by combining together the conditional local volatility process, the volatility regime process and the

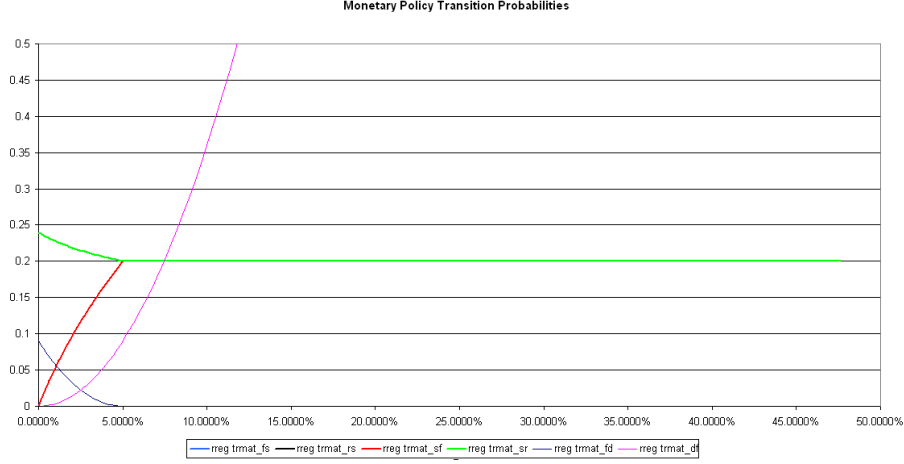


Figure 4: Transition probability intensities between different drift regimes, as a function of the value of the underlying rate L

drift regime process. Let $\mathcal{L}_{\Omega_L}^j$ be the Markov generator for the subordinated conditional local volatility process defined on the lattice Ω_L , \mathcal{L}_{Ω_a} the Markov generator for the stochastic volatility regime process defined on the lattice Ω_a and \mathcal{L}_{Ω_b} the Markov generator for the stochastic drift regime process defined on the lattice Ω_b . Assuming no cross terms, the Markov generator for the combined process is defined on a finite dimensional lattice $\Omega = \Omega_L \times \Omega_a \times \Omega_b$ and is simply given by:

$$\mathcal{L}_{\Omega}(x_1, a_1, b_1; x_2, a_2, b_2) = \mathcal{L}_{\Omega_L}^j(x_1, x_2)\delta_{a_1 a_2}\delta_{b_1 b_2} + \quad (36)$$

$$\mathcal{L}_{\Omega_a}(a_1, a_2)\delta_{x_1 x_2}\delta_{b_1 b_2} + \quad (37)$$

$$\mathcal{L}_{\Omega_b}(b_1, b_2)\delta_{x_1 x_2}\delta_{a_1 a_2}$$

where δ is the Kronecker delta and $(x_1, x_2), (a_1, a_2), (b_1, b_2)$ are two generic states for the conditional local volatility process, the volatility regime process and the drift regime process respectively.

In a more compact form we write $\mathcal{L}_{\Omega}(y_1; y_2) = \mathcal{L}_{\Omega}(x_1, a_1, b_1; x_2, a_2, b_2)$, where we remind that in our notation $y_i = (x_i, a_i, b_i)$. To gain an intuition on the meaning of the matrix element $\mathcal{L}_{\Omega}(y_1; y_2)$ we note that, for each element off the main diagonal, $\mathcal{L}_{\Omega}(y_1; y_2)dt$ represents the transition probability to move from a state $L(x_1)$ in the b_1 volatility regime and a_1 drift regime to a state $L(x_2)$ in the b_2 volatility regime and a_2 drift regime in an infinitesimal time dt . We are therefore describing a three-dimensional problem: one for the underlying LIBOR rate, one for the stochastic volatility regime and one for the stochastic drift regime.

The Markov propagator of the combined process is computed by taking the exponential of the combined Markov generator:

$$U(y_i, t_i; y_{i+1}, t_{i+1}) = e^{\mathcal{L}\Omega(t_{i+1}-t_i)}(y_i, y_{i+1}). \quad (38)$$

Since our underlier is a 3-month rolling LIBOR rate though, we are not interested in the pricing kernel but rather in the discounted transition probability kernel, which in the risk neutral measure, is given by:

$$G(y_i, t_i; y_j, t_j) = \sum_{y_{i+1}, y_{i+2}, \dots, y_{j-1}} \prod_{k=i+1}^j U(y_{k-1}, t_{k-1}; y_k, t_k) \frac{1}{1 + \tau L_{t_{k-1}}} \quad (39)$$

The 3-month discounted transition probability kernel is obtained by discounting the 3-month propagator $U(y_{k-1}, t_{k-1}; y_k, t_k)$ between times $(t_{k-1}; t_k)$ with the 3-month discount factor $\frac{1}{1 + \tau L_{t_{k-1}}}$. In order to obtain the discounted transition probability kernel between two generic times $(t_i; t_j)$, with $t_j - t_i > 0.25$, the elementary discounted kernel is compounded $n = \frac{t_j - t_i}{0.25}$ times summing across all intermediate states.

4 Calibration

The model is calibrated to the discount curve and the so called swaption volatility cube. The model implied terminal correlation structure is also observed and compared with historical data.

A key driver in the calibration exercise has been to find a model parameterisation that not only explains the market at one particular date, but such that the yield curve shapes obtained as a function of the initial condition are consistent with historical data. This is instrumental in order to aim for a stable model parameterisation. The calibration is first performed with a time homogeneous parameterisation, and subsequently refined with minor deterministic time dependent shifts for a perfect match to the term structure of the interest rates.

In the time homogeneous case, discount factors conditional on an initial state y_i are given by:

$$Z(y_i; t_i, t_j) = \sum_{y_j} G(y_i, t_i; y_j, t_j) \quad (40)$$

and swap rates can be evaluated as follows:

$$SR(y_i; t_i, t_j) = \frac{1 - Z(y_i; t_i, t_j)}{BPV(y_i; t_i, t_j)} \quad (41)$$

where $BPV(y_i; t_i, t_j) = \sum_{k=i+1}^j \tau Z(y_i; t_i, t_k)$. For the calibration to the swaption volatility cube, one needs to evaluate the price of a portfolio of European swaptions, for different maturities, underlying tenors and strikes. In particular,

the value of an European swaption maturing at time $t_j > t_i$, where the maturity of the underlying swap is $t_k > t_j$ and struck at k is given by, in the risk neutral measure:

$$SO(y_i; t_i, t_j) = \sum_{y_j} G(y_i, t_i; y_j, t_j) (SR(y_j; t_j, t_k) - k)_+ BPV(y_j; t_j, t_k). \quad (42)$$

Time inhomogeneities can be introduced in the model in order to obtain a perfect fit to the term structure of interest rates. This can be achieved by normalising the discounted transition probability with a deterministic time dependent factor:

$$G(y_i, t_i; y_j, t_j) \rightarrow \bar{G}(y_i, t_i; y_j, t_j) = G(y_i, t_i; y_j, t_j) \frac{Z_{mkt}(t_j)Z(\bar{y}_0; t_0, t_i)}{Z_{mkt}(t_i)Z(\bar{y}_0; t_0, t_j)} \quad (43)$$


where $Z_{mkt}(t_j)$ is the market discount factor with maturity t_j , \bar{y}_0 is the spot value at current time t_0 of the underlying lattice process. If the time homogeneous parameterisation is well calibrated, these time dependent change will produce shifts in the interest rates which are typically around 10 basis points.

Fig. (5) reports all possible shapes of yield curve obtained with the model, as a function of the initial point x_i and conditional on each drift regime. One can notice that the shape of the discount curve is only marginally affected by the instantaneous volatility and is largely influenced by the prevailing monetary regime. This is not surprising as the volatility process mean reverts of time scales much shorter than the time scales for the monetary policy dynamics. Hence, volatility of long term rates, spread between rates and the correlation between forward rates across the term structure are mainly controlled by the stochastic drift term.


The implied volatility of swaptions at various maturities is characterised by a persistent skew. The model is capable of describing this persistent skew effect, and the implied volatility smiles are qualitatively in line with the market, as reported in fig. (6).

An interesting observation allows to gain more insight on the impact of the drift regime and in particular the deflationary regime. In fact, by removing the deflationary regime, in our experience, it has not been possible to calibrate the low strikes portion of the smile for long dated maturities without affecting the dynamics of the process at earlier maturities. Fig. (7) shows that the 20y into 30y swaption smile at low strikes is too flat compared to the market. The introduction of a deflationary regime allows to increase the probability of the rate of spanning regions of low values at long maturities, and therefore lifts up the smile at low strikes, without considerably affecting the smile at short and medium maturities. This example highlights the important fact that the monetary policy process can be used to provide direct control to the long dated behaviour of the rate process.


Our calibration process has focused primarily on a qualitative fit of the swaption volatility cube, however we also note that the model implied short




(a) Deflation



(b) Falling



(c) Stable



(d) Rising

Figure 5: Yield curve shapes as a function of the initial point and conditional on the drift regime

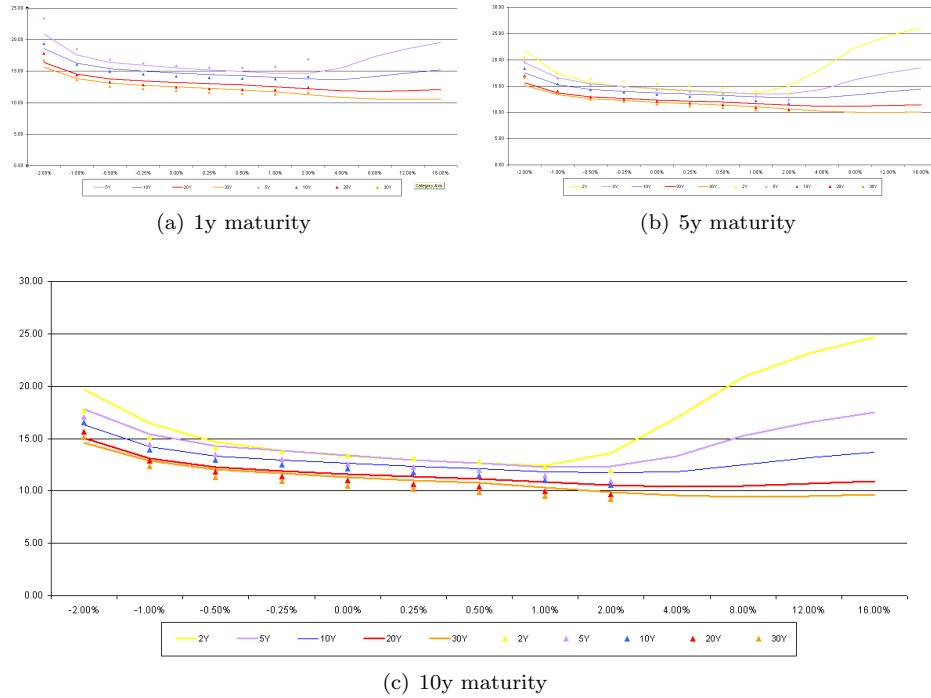


Figure 6: Model (lines) and market (dots) swaption volatility smile at different maturities and for different tenors




Figure 7: Model (lines) and market (dots) swaption volatility smile at 20y maturity with and without deflation

term correlation also fall in line with historical estimates across tenors and maturities.

Tenors	vol	weekly model correlations						
		0.25	0.5	1	2	5	10	20
0.25	16.27	100.00	99.85	98.79	94.60	78.53	61.66	50.07
0.5	15.73	99.85	100.00	99.48	96.21	81.73	65.77	54.57
1	15.18	98.79	99.48	100.00	98.48	87.14	73.02	62.64
2	14.49	94.60	96.21	98.48	100.00	94.32	83.69	75.00
5	13.92	78.53	81.73	87.14	94.32	100.00	97.09	92.48
10	13.35	61.66	65.77	73.02	83.69	97.09	100.00	98.83
20	11.77	50.07	54.57	62.64	75.00	92.48	98.83	100.00

Figure 8: One week term correlation structure

5 Applications

We have applied the stochastic monetary policy model to callable swaps and callable CMS spread range accruals. We have priced a 20y maturity payer callable swap, where the counterparty paying a stream of fixed rate payments and receiving a Libor leg has the option to cancel the swap every year. For the callable CMS spread range accrual we have considered the case where the counterparty paying the range accrual leg and receiving the Libor leg has the option to cancel the swap every year. In our example, the range accrual coupons are contingent on the spread between the 10y and 2y CMS rates. The range accrual coupon c_i for each coupon period i is calculated as follows:

$$c_i = k \frac{n_i}{N_i} = \frac{\sum_{j_i=1}^{N_i} \mathbf{1}_{\{SR(t_j, t_j+10y) - SR(t_j, t_j+2y) > b\}}}{N_i} \quad (44)$$

where: n_i is the total number of days when the 10y-2y CMS spread sets above a low barrier level in the coupon period i and N_i is the total number of observations in the same coupon period i .

In order to appreciate the impact of the monetary policy implied by the model, it is instructive to analyse how the price changes as a function of the initial condition on the spot rate and conditional on the drift regime.

5.1 Callable swaps

For a given initial condition, the price of a callable swap is always an increasing function in the monetary policy regime: the lowest price being obtained in the deflationary regime and the highest price in the rising regime. This result is not surprising and also quite obvious: the value of a callable payer swap is expected

to be higher in a rising regimes, when rates are biased on a rising trend, than in a falling regime, when interest rates are biased on a falling trend. Fig (9) show the price functions, as a function of the initial condition on the rate, and the exercise boundaries as a function of time, for a 20y maturity callable swap struck at 5%.





Figure 9: Price functions of a 20y maturity callable swap, as a function of the initial condition on the rate, and exercise boundaries as function of time

5.2 Callable CMS spread range accruals


For the case of callable CMS spread range accruals, the analysis requires more care. In fact, in this case, the monetary policy regime has an impact on both the level of the rates and the level of the CMS spread. In this example, it becomes apparent the effect of the correlation framework based on the dynamic conditioning to monetary policy regimes. In addition to the price functions and exercise boundary graphs, shown in fig.(10), it is instructive to analyse the range accrual probabilities, as a function of the initial condition on the rate, fig.(11). These probabilities indicate the average probability, within a coupon period, that the CMS spread sets above the specified barrier level, which in our example was set to 10 basis points. One can note that, for the case of the deflationary regime, the model generate nearly zero range accrual probabilities for most initial conditions, a part from extreme low initial condition levels. This means that, conditional on a deflationary regime, the model implies a flat or inverted yield curve. This is what one would expect taking into account our the definition of the deflationary regime: interest rates are constrained to stay at very low level once these levels have been reached.

6 Conclusions

We have presented an interest rate model which incorporates stochastic monetary policy in addition to a conditional local volatility component, jumps and



(a) Price Functions



(b) Exercise Boundaries

Figure 10: Price functions of a 20y maturity callable CMS spread range accrual, as a function of the initial condition on the rate, and exercise boundaries as function of time

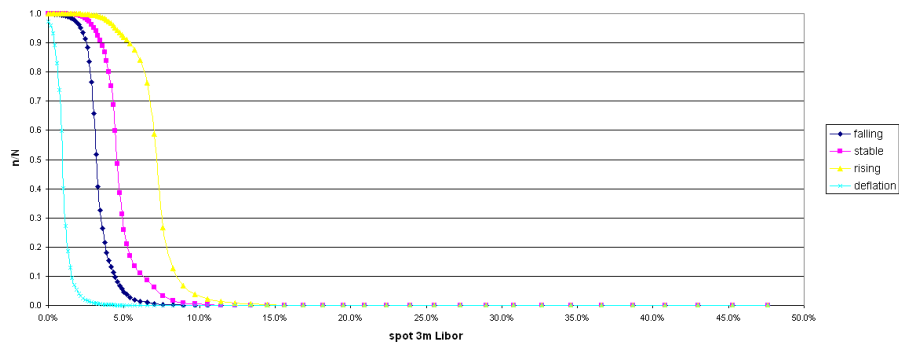


Figure 11: Range accrual probabilities conditional on drift regimes

stochastic volatility. This is a novel approach in interest rate modelling and it represents the main contribution of this paper. Time-homogeneity has been a key driver in the construction of our model and we show that, with the additional degree of freedom provided by the drift process, the model is able to describe a wide variety of yield curve shapes. As a result, we find that the monetary policy component helps achieving model consistency and allows to obtain a good qualitative fit to the swaption volatility cube and the correlation structure.

The model is built with Markov chains on finitely many states and is solved by means of continuous-time lattices. The mathematical framework and the numerical algorithm for the construction of a continuous-time lattice has been described in this work and this represents a refinement to what we have proposed in our previous works.

Finally, we have shown an application of the model to callable swaps and callable CMS spread range accruals. The dependence of the price functions and the range accrual probabilities to the drift regimes reveals the impact of stochastic monetary policy in interest rate derivative pricing. The modelling framework proposed in this paper allows to directly specify views on economic scenarios and analyse the impact of these views on the valuation and risk management of exotic interest rate structures.

References

- Albanese, C. and A. Mijatovic (2006), Convergence rates for diffusion on continuous-time lattices. Working paper.
- Albanese, Claudio and Alexey Kuznetsov (2003), Discretization schemes for subordinated processes. *Mathematical Finance*. To appear.
- Albanese, Claudio and Alexey Kuznetsov (2005), ‘Affine lattice models’, *International Journal of Theoretical and Applied Finance* **8(2)**, 223–238.
- Albanese, Claudio and Manlio Trovato (2005), A stochastic volatility model for bermuda swaption and callable cms swaps. Working paper.
- Bochner, S. (1955), *Harmonic analysis and the theory of probability*, University of California Press.
- Karatzas, I. and S. Shreve (1991), *Brownian Motion and Stochastic Calculus*, Springer.
- Kushner, H. and P. Dupuis (2001), *Numerical Methods for Stochastic Control Problems in Continuous Time*, Springer.
- Moler, C. and C.V. Loan (2003), ‘Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later’, *SIAM Review* **45(1)**, 3–30.

Rebonato, Riccardo and Mark Joshi (2001), A joint empirical and theoretical investigation of the modes of deformation of swaption matrices: implications for model choice. QUARC Working paper.

Threfethen, Lloyd N. (1999), 'Computation of pseudospectra', *Acta Numerica*, Cambridge University Press pp. 247–295.