Monetary Policy Risk and CMS Spreads *

Claudio Albanese † Manlio Trovato ‡

May 12, 2007

Abstract

Central banks' monetary policies are regarded by financial institutions as a key driver for the definition of their interest rate hedging strategies. However this valuable information is not directly incorporated in most derivative pricing models commonly used in financial institutions. We present a novel approach to interest rate modelling, which incorporates a direct specification of stochastic monetary policy within an arbitrage-free context. In particular, we take the 3-month spot LIBOR rate as modelling primitive and propose a three factor interest rate term structure model solved on a continuous-time lattice. The model is constructed with local volatility, stochastic volatility regimes and stochastic drift regimes and can be formally written as $dL_t = \mu_{a_t}(L_t)dt + \sigma_{b_t}(L_t)dW_t$. The drift and the volatility terms are stochastic and driven by the processes a_t and b_t, which can be made correlated to the rates themselves.

We show that, with a nearly time-homogeneous parameterisation, the model can achieve a persistent smile structure across maturities, in agreement with the EUR market, and that the model implied correlation structure is consistent with historical estimates. We apply the model to callable swaps and callable CMS spread range accruals and analyze the impact of monetary policy to the valuation of these derivative contracts.

1 Introduction

The prevailing monetary policy regime and its future market expectation have a fundamental impact on interest rates trading activity. Central banks' monetary policies are actively monitored by financial institutions and are key for the definition of their interest rate hedging strategies. It has been shown that the market expectation on monetary policy is well reflected in plain vanilla options prices (Söderlind, 2000). In particular, futures, future options, foreign exchange options and swaps on the overnight rate define a clear picture of the expected monetary policy up to 36 months forward. However this valuable information

*The authors would like to thank Gianni Adragna for his useful comments.

†Claudio Albanese. E-mail: claudio@level3finance.com

‡Manlio Trovato, Merrill Lynch and Imperial College. E-mail: manlio_trovato@ml.com

is not directly incorporated in most derivative pricing models commonly used by financial institutions for the valuation and risk management of long dated interest rate exotics.

The difficulty in evaluating the impact of monetary policy for complex derivative contracts, like callable CMS spread range accruals, lies in the development of a tractable and coherent modelling framework. In this work we present a novel approach to interest rate modelling, which incorporates the direct specification of the stochastic monetary policy process. We take the 3-month LIBOR spot rate L_t as our modelling primitive. As this is not an asset price process, the drift term is not constrained by the no-arbitrage condition and can be directly specified.

We propose an arbitrage-free interest rate model built upon the specification of a conditional local volatility process, a volatility regime process and a drift regime process and enforce a nearly time-homogeneous parameterisation. The process for L_t can formally be written as follows:

$$dL_t = \mu_{a_t}(L_t)dt + \sigma_{b_t}(L_t)dW_t \tag{1}$$

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 which can be made correlated to the rates themselves. The direct specification of the drift functions μ and the drift process b_t provides a natural framework for imposing one's view on the economic outlook into the model. The remaining degrees of freedom are used to make the model consistent with the swaptions volatility cube and historical estimates of the correlation structure.

Our model specification is entirely non-parametric and a noiseless numerical framework is obtained using fast exponentiation and the theory of Abelian processes. See (Albanese and Trovato, 2005) and references therein for a more technical discussion. Since the application domain is not limited by analytical tractability, the rate process used for pricing resembles quite closely the historical process.

In conclusion, the main contribution of this article is the introduction of stochastic monetary policy in interest rate modeling. The use of operator methods and continuous-time lattices also represents a distinctive feature of our modeling framework.

The sequel is organized as follows: we first discuss the basic concepts for the construction of continuous-time lattices; we then describe the specification of the interest rate model and its calibration. Finally, we present an application to callable swaps and callable CMS spread range accruals.

2 Continuos-time lattices

The basic concepts underlying the construction of continuous-time lattices are best explained for the simple case of diffusion processes. In this case, we consider a generic asset price process F_t and assume that it satisfies the following SDE:

$$dF_t = \mu(F_t)dt + \sigma(F_t)dW_t. \tag{2}$$

Let $\mathcal{L}$ be the Markov generator for F_t, defined on a continuum state space. As a function of the backward variables f and t, the probability density function $u(f,t;F,T)$ is known to satisfy the backward Kolmogorov equation (Karatzas and Shreve, 1991):

$$\frac{\partial u}{\partial t} + \mathcal{L}u = 0 \tag{3}$$

with terminal condition $u(f,T;F,T) = \delta(f-F)$. The solution to eq. (3) can be formally written as:

$$u(f,t;F,T) = e^{\mathcal{L}(T-t)}(f,F). \tag{4}$$

where in this notation we mean that the operator $e^{\mathcal{L}(T-t)}$ is applied to the variables (f,F). It is therefore obvious that, in order to solve for the probability density function u all we need is the specification of the Markov generator $\mathcal{L}$. It can be shown (see (Albanese and Mijatovic, 2006) and (Kushner and Dupuis, 2001)) that an appropriate Markov chain approximation of the diffusion process (2) can be constructed by imposing the following local moment matching conditions:

$$\mathbb{E}_t\left[F_{t+dt} - F_t\right] = \mu dt \tag{5}$$

$$\mathbb{E}_t\left[(F_{t+dt} - F_t)^2 - \mathbb{E}_t\left[F_{t+dt} - F_t\right]^2\right] = \sigma^2 dt \tag{6}$$

Let $\mathcal{L}_\Omega$ be the following finite difference discretization of the operator $\mathcal{L}$, defined on a finite state space $\Omega = \{0,1,2,...,N\}$:

$$\mathcal{L}_\Omega := \mu\nabla_h + \frac{\sigma^2}{2}\Delta_h \tag{7}$$

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 moment matching conditions in terms of the discretised Markov generator $\mathcal{L}_\Omega$ as follows:

$$\sum_{y\in\Omega} \mathcal{L}_\Omega(x,y)(F(y) - F(x)) = \mu(F(x)) \tag{8}$$

$$\sum_{y\in\Omega} \mathcal{L}_\Omega(x,y)(F(y) - F(x))^2 = \sigma^2(F(x)). \tag{9}$$

for all $x \in \Omega$. The above two equations, together with the probability conservation condition:

$$\sum_{y\in\Omega} \mathcal{L}_\Omega(x,y) = 0 \tag{10}$$

uniquely identify the Markov generator for the diffusion process (2).

Once the Markov generator for the approximating Markov chain has been constructed, the discretised transition probability kernel u_Ω can be simply obtained by computing the exponential of the Markov generator matrix itself:

$$u_\Omega(x,t;y,T) = e^{\mathcal{L}_\Omega(T-t)}(x,y). \tag{11}$$

This can be efficiently calculated with the *scaling and squaring* method (Moler and Loan, 2003) by means of Level-3 BLAS numerical routines: in this case matrix-matrix multiplication for medium to large matrices of dimension n appears to scale as $n^{2+\varepsilon}$, which is a substantial improvement to the naive triple loop implementation that scales as n^3. Moreover, by capturing the numerical bottleneck in matrix-matrix multiplications one obtains a framework which is ideally suited for hardware acceleration by means of the emerging GPU co-processors.

2.1 The stochastic monetary policy model

We introduce 70 states for the conditional local volatility process, 3 states for the volatility regime variable and 4 states for the drift regime variable. Let the state space for each factor be Ω_L=$\{0,1,...,69\}$, Ω_a=$\{0,1,2\}$ and Ω_b=$\{0,1,2,3\}$. 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 $b_i \in \Omega_b$.

2.1.1 The conditional local volatility processes

Conditional to a volatility regime state a and a drift regime state b, we define an approximating Markov chain associated to the diffusion process for the 3-month LIBOR spot rate. The Markov generator $\mathcal{L}_{\Omega_L}$ is constructed by imposing probability conservation and the local consistency properties, in the sense of equation (10), (8) and (9) respectively.

Drift specification The drift functions can be specified consistently with historical estimates and one's view on the economic outlook. The estimation of the drift functions from historical time series is a very active research field in econometric. (Andersen et al., 2004) consider parametric jump diffusions with stochastic volatility, whereas (Stanton, 1997) discusses a fully non parametric approach for one factor diffusion processes and show that the drift is substantially non linear. In particular, they show that the mean reversion effect appears to be negligible at low to medium level of rates, and that it rapidly increases at high rates. In our model, the drift functions are conditional to the drift regime variable and are constructed with a series of piecewise linear functions of the underlying rate (fig. 1(a)). In our parametrization choice, for the `stable` and `rising` regimes, the drift is constant at low rates, whereas it has a mean reverting behavior at medium to high rates, consistently with (Stanton, 1997). For `falling` and `deflation` regimes, the higher the level of rates, the more the process is pulled towards low levels: this makes intuitive sense as, by definition, in these regimes the rate process is biased towards low interest rate levels.

Local volatility specification Whilst the construction of the drift functions is primarily based on economically meaningful views, the local volatility specification aims to achieve consistency with the observed skew in the swaption market. The local volatility is of a *CEV* type (fig. (1(b))), which has received much attention in finance:

$$\sigma_b(L(x)) = \overline{\sigma}_b L(x)^{\beta(L(x))}. \tag{12}$$

The exponent is made dependent of the level of the underlying rates:

$$\beta(L(x)) = \frac{A + L(x)}{B + L(x)} \tag{13}$$

where A and B are solved by imposing $\beta(0) = \beta_0$ and $\beta(L_{ref}) = \beta_{ref}$ (fig. (2)). Such state dependent exponent can be interpreted as a state dependent distribution assumption: interest rates behave close to normal at very low levels and close to lognormal at high levels.

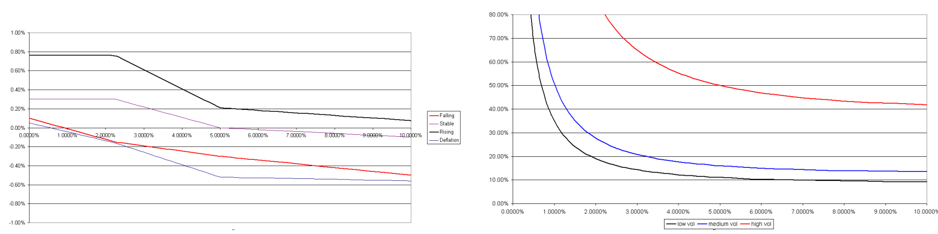

(a) Drift

(b) Lognormal local vol

Figure 1: Specification of drift and local volatility, conditioned on the drift and volatility regimes respectively

2.1.2 Modelling the dynamics of stochastic volatility regimes

The process for the stochastic volatility is not easily observed in interest rates. Payoffs whose primitive underlier is the volatility itself are still scarce and illiquid, although they have been receiving increased attention in very recent times. (Rebonato and Joshi, 2001) argue that the volatility process experiences patterns which can be effectively described by means of regimes. Consistently with

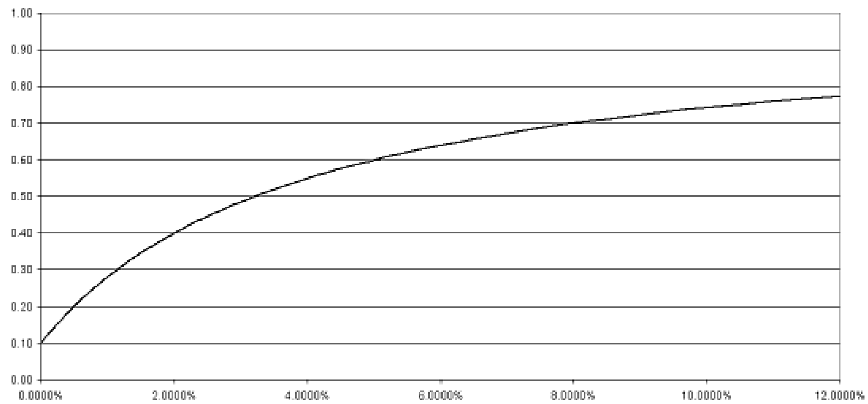

Figure 2: Specification of β as a function of the LIBOR rate

this view, the volatility process in our model is allowed to switch between three possible states: `low`, `medium` and `high`. 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 behavior 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. In our model parameterizations, the Markov generator $\mathcal{L}_{\Omega_a}$ is set to be independent of the rates and we have used the following values, which allow to achieve a good degree of consistency with the swaption volatility cube:

$$
\begin{array}{cccccc}
 & \mathit{high} & \mathit{med} & \mathit{low} & & \\
 & \downarrow & \downarrow & \downarrow & & \\
\mathcal{L}_{\Omega_a} = & \left(\begin{array}{c} -2.5 \\ 0.2 \\ 0 \end{array}\right. & \begin{array}{c} 2.5 \\ -0.6 \\ 0.4 \end{array} & \left.\begin{array}{c} 0 \\ 0.4 \\ -0.4 \end{array}\right) & \begin{array}{c} \leftarrow \\ \leftarrow \\ \leftarrow \end{array} & \begin{array}{c} \mathit{high} \\ \mathit{medium} \\ \mathit{low} \end{array}
\end{array}
$$

2.1.3 Modelling the dynamics of a stochastic monetary policy

Historical time series analysis We analyse the historical time series for 3-month LIBOR spot rate for the three major currencies: EUR, USD and JPY (fig.(3)). It is obvious that the process is driven by trends which last for periods that may 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, respectively. 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 `deflation` drift regime which is associated to this phenomenon.

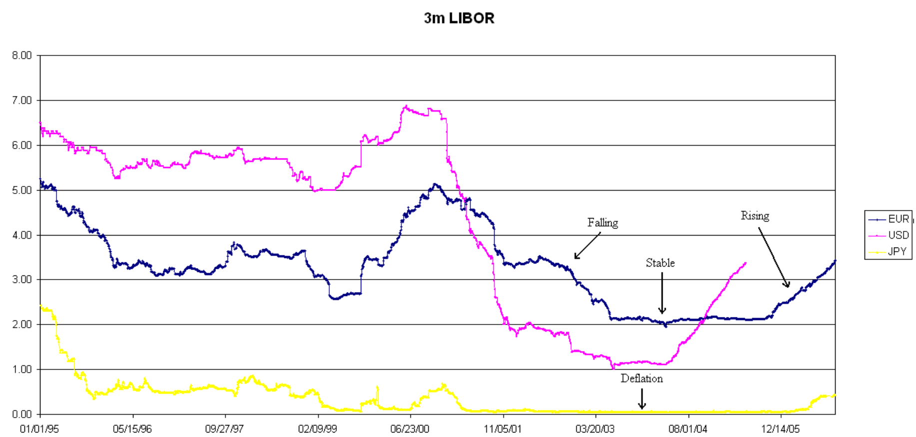

Figure 3: spot 3m Libor rate historical time series

Construction of the Markov generator The elements of the Markov generator $\mathcal{L}_{\Omega_b}$ for 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. In particular, the following observations have driven our parameterizations choice:

- 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;
- 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; the inverse applies for `deflation` to `falling`;
- the probability of switching from a `stable` to a `falling` regime, and from a `rising` to a `stable` regime, increases up to a constant as rates increase

up to a reference level; the inverse applies for `falling` to `stable` and `stable` to `rising`;

One should note that the above behavior 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. Our model provides a natural mathematical framework for such specification.

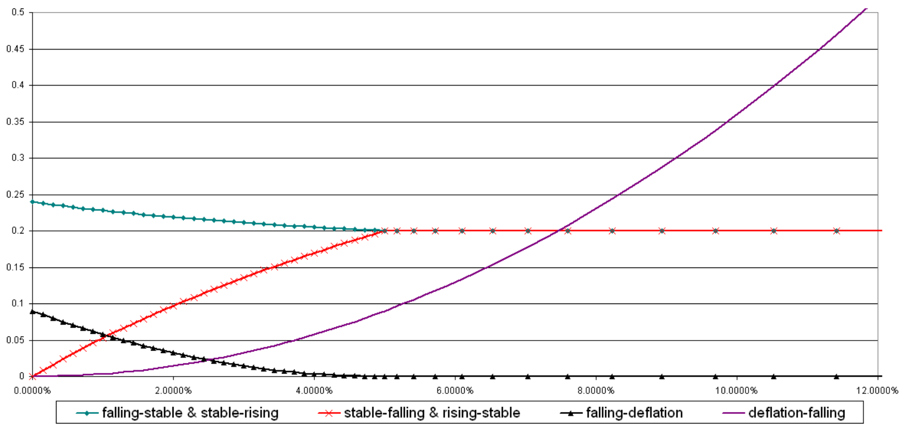

Figure 4: Transition probability intensities between different drift regimes, as a function of the value of the underlying rate L

2.1.4 The combined process for the spot LIBOR rate

We complete the construction of our interest rate model by combining the three factors together. Assuming no cross terms, the Markov generator for the combined process, defined on a finite dimensional lattice $\Omega = \Omega_L \times \Omega_a \times \Omega_b$, is simply given by:

$$\begin{aligned} \mathcal{L}_\Omega(x_1, a_1, b_1; x_2, a_2, b_2) = \; & \mathcal{L}^j_{\Omega_L}(x_1, x_2)\delta_{a_1 a_2}\delta_{b_1 b_2} + \\ & \mathcal{L}_{\Omega_a}(a_1, a_2)\delta_{x_1 x_2}\delta_{b_1 b_2} + \\ & \mathcal{L}_{\Omega_b}(b_1, b_2)\delta_{x_1 x_2}\delta_{a_1 a_2} \end{aligned} \tag{14}$$

where δ is the Kronecker delta.

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 $y_i = (x_i, a_i, b_i)$. The Markov propagator of the combined process is computed by taking the exponential of the combined Markov generator:

$$U(y_i, t_i; y_{i+1}, t_{i+1}) = e^{\mathcal{L}_\Omega (t_{i+1}-t_i)}(y_i, y_{i+1}). \tag{15}$$

In the spot measure with simple compounding we use, the discounted transition probability kernel is given by:

$$G(y_i, t_i; y_j, t_j) = \sum_{y_{i+1}, y_{i+2}, \ldots, y_{j-1}} \prod_{k=i+1}^{j} U(y_{k-1}, t_{k-1}; y_k, t_k) \frac{1}{1 + \tau L_{t_{k-1}}} \tag{16}$$

Therefore $G(y_{k-1}, t_{k-1}; y_k, t_k)$ is obtained by discounting the 3-month propagator $U(y_{k-1}, t_{k-1}; y_k, t_k)$ with the 3-month discount factor $\frac{1}{1+\tau L_{t_{k-1}}}$. For 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.

3 Calibration

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

The calibration is first performed with a time-homogeneous parameterizations and subsequently refined with minor deterministic time dependent lattice dilations to obtain a perfect match to the current term structure of the interest rates. For model calibration, discount factors, swap rates and swaption prices, conditional on an initial state y_i, must be evaluated on the lattice. In the time homogeneous case, these are given by:

$$\begin{aligned} Z(y_i; t_i, t_j) &= \sum_{y_j} G(y_i, t_i; y_j, t_j) \qquad (17) \\ SR(y_i; t_i, t_j) &= \frac{1 - Z(y_i; t_i, t_j)}{BPV(y_i; t_i, t_j)} \\ SO(y_i; t_i, t_j, t_k) &= \sum_{y_j} G(y_i, t_i; y_j, t_j) \left(SR(y_j; t_j, t_k) - k \right)_+ BPV(y_j; t_j, t_k) \end{aligned}$$

where $BPV(y_i; t_i, t_j) = \sum_{k=i+1}^{j} \tau Z(y_i; t_i, t_k)$ and $SO(y_i; t_i, t_j, t_k)$ is the value of an European swaption, with maturity $t_j > t_i$, underlying swap maturity $t_k > t_j$ and strike k.

A key driver in the calibration exercise has been to find a model parameterizations 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. 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, and in particular its long term behavior, is largely influenced by the prevailing monetary regime. 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.

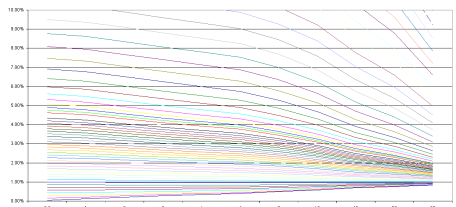

(a) Deflation

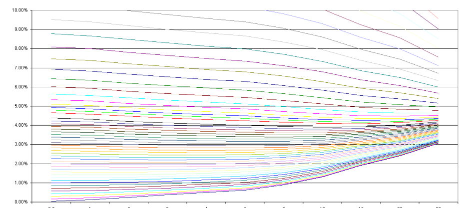

(b) Falling

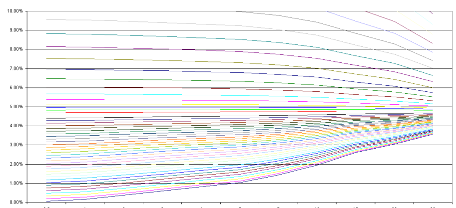

(c) Stable

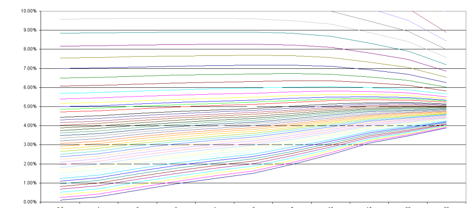

(d) Rising

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

The implied volatility of swaptions at various maturities is characterized by a persistent skew. The model is capable of describing this effect, and the implied volatility smiles are qualitatively in line with the market (fig. (6)). This effect is primarily generated by the local volatility and stochastic volatility components at short to medium maturities. However, at longer maturities, stochastic drift becomes the predominant factor. As an example, we show that the low strikes smile at 20y maturity is mainly explained by the deflation scenario (fig. (7)). This makes intuitive sense, as the deflation scenario increases the probability of the rate to span and stay at very low levels.

Finally, we note that the model implied short term correlations are also in line with historical estimates across tenors and maturities (fig. (8)). In this table, the 1 week correlation, evaluated at t_i, between two swap rates of underlying tenors τ_1 and τ_2 is calculated as follows:

$$\rho_\delta\left(t_i;\tau_1,\tau_2\right)=\frac{\mathbb{E}_{t_i}\left[\left(S^{\tau_1}_{t_i+\delta}-S^{\tau_1}_{t_i}\right)\left(S^{\tau_2}_{t_i+\delta}-S^{\tau_2}_{t_i}\right)\right]}{\sqrt{\mathbb{E}_{t_i}\left[\left(S^{\tau_1}_{t_i+\delta}-S^{\tau_1}_{t_i}\right)^2\right]}\sqrt{\mathbb{E}_{t_i}\left[\left(S^{\tau_2}_{t_i+\delta}-S^{\tau_1}_{t_i}\right)^2\right]}} \tag{18}$$

where $\delta=1/52$ and we have used the notation $S^{\tau}_{t_i}=SR(t_i,t_i+\tau)$.

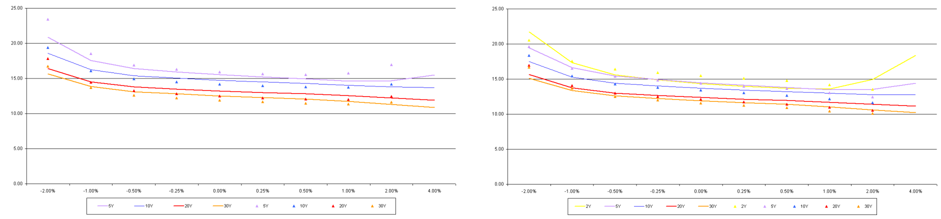

(a) 1y maturity (b) 5y maturity

(c) 10y maturity (d) 20y maturity

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

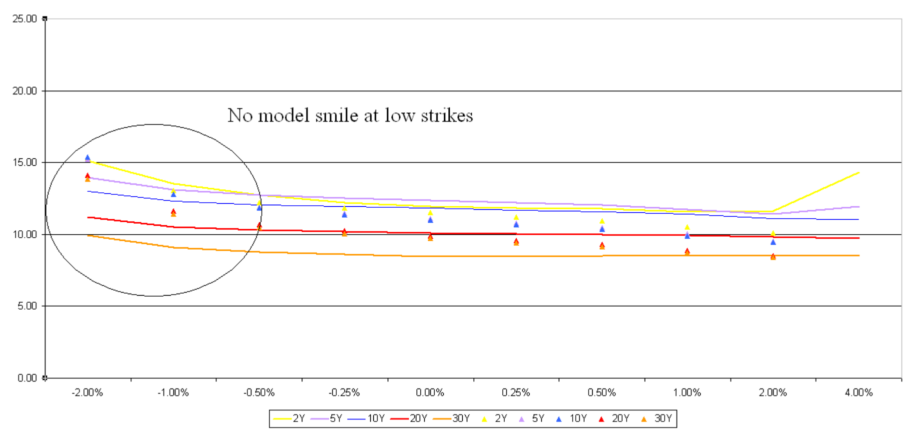

Figure 7: Model (lines) and market (dots) swaption volatility smile at 20y maturity without deflation regime

Tenors	0.25	0.5	1	2	10	20	30
0.25	100.00	99.85	98.74	94.38	60.30	62.09	50.60
0.5	99.85	100.00	99.46	96.06	64.59	66.15	55.04
1	98.74	99.46	100.00	98.43	72.15	73.31	63.02
2	94.38	96.06	98.43	100.00	83.25	83.87	75.24
10	60.30	64.59	72.15	83.25	100.00	97.12	92.54
20	62.09	66.15	73.31	83.87	97.12	100.00	98.84
30	50.60	55.04	63.02	75.24	92.54	98.84	100.00

weekly model correlations

Figure 8: One week term correlation structure

4 Applications

4.1 Callable swaps

We plot the price functions and exercise boundaries for a 20y maturity callable payer swap struck at 5% with annual exercise schedule (fig (9)). For a given initial condition, the price of a callable payer 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 regime, when rates are biased on a rising trend, than in a falling regime, when interest rates are biased on a falling trend.

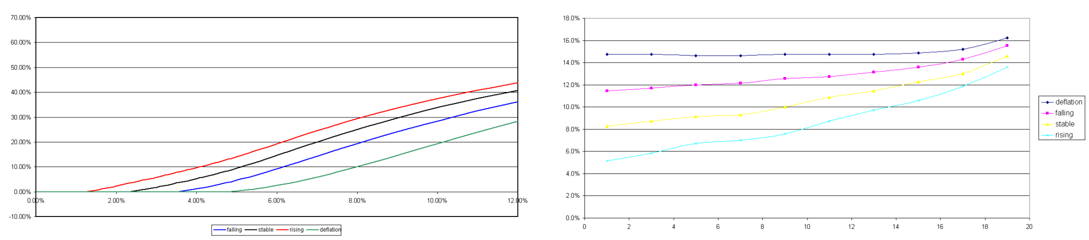

(a) Price Functions (b) Exercise Boundaries

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

4.2 Callable CMS spread range accruals

We price a callable CMS spread range accrual, 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 c_i are contingent on the spread between the 10y and 2y CMS rates and, for each coupon period i, are calculated as follows:

$$c_i = k\frac{n_i}{N_i} = \frac{\sum_{j=1}^{N_i} \mathbf{1}_{\{SR(t_j,t_j+10y)-SR(t_j,t_j+2y)>b\}}}{N_i} \tag{19}$$

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

In this case, the analysis requires more care. In fact, the monetary policy regime has 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. We plot the price functions and the exercise boundaries in fig.(10) and the range accrual probabilities in fig.(11). The latter reports the average probabilities, within a coupon period and conditional on the state at the start of the period itself, that the CMS spread sets above the specified barrier level, which in our example was set to 10 basis points. In a time-homogeneous parameterisation, these are the same for every coupon period. 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 very low initial condition levels. This is consistent with the fact that, conditional on a deflationary regime, the model implies a flat or inverted yield curve.

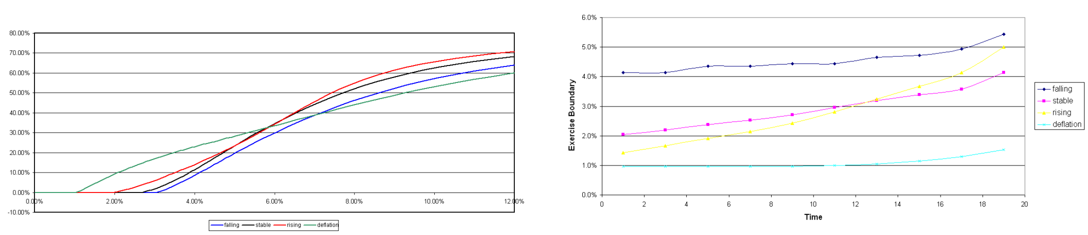

(a) Price Functions

(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

5 Conclusions

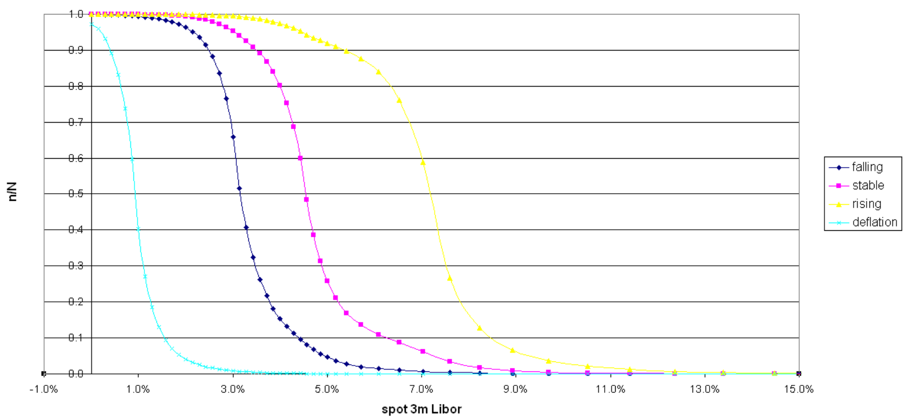

Figure 11: Range accrual probabilities conditional on drift regimes, as a function on the initial condition on the rate

We have presented a novel approach to interest rate modelling, which incorporates a direct specification of stochastic monetary policy in addition to a conditional local volatility component and stochastic volatility. The model is built with Markov chains on finitely many states and solved by means of continuous-time lattices, which have also been briefly discussed in this work. We find that the additional degree of freedom provided by the drift process helps achieving model consistency and allows to obtain a good qualitative fit to the swaption volatility cube and the correlation structure.

We have applied the model to callable swaps and callable CMS spreads and have shown that, in our modelling framework, it is possible to evaluate the impact of the monetary policy on the valuation of these interest rate derivative contracts.

References

Albanese, Claudio and Aleksandar Mijatovic (2006), Convergence rates for diffusion on continuous-time lattices. Working paper.

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.

Andersen, Torben G., Luca Benzoni and Jesper Lund (2004), Stochastic volatility, mean drift, and jumps in the short-term interest rate. Working paper.

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.

Söderlind, Paul (2000), 'Market expectations in the uk before and after the erm crisis', *Economica* **67**, 1–18.

Stanton, Richard H. (1997), 'A nonparametric model of term structure dynamics and the market price of interest rate risk', *Journal of Finance* **52**, 1973–2002.