INTEREST RATES AND FX MODELS

6. LIBOR Market Model

Andrew Lesniewski

Courant Institute of Mathematical Sciences
New York University
New York

March 6, 2013

Contents

1 Introduction

The real challenge in modeling interest rates is the existence of a term structure of interest rates embodied in the shape of the forward curve. Fixed income instruments typically depend on a segment of the forward curve rather than a single point. Pricing such instruments requires thus a model describing a stochastic time evolution of the entire forward curve.

There exists a large number of term structure models based on different choices of state variables parameterizing the curve, number of dynamic factors, volatility smile characteristics, etc. The industry standard for interest rates modeling that has emerged over the past few years is the LIBOR market model (LMM). Unlike the older approaches (short rate models which we do not discuss in these lectures), where the underlying state variable is an unobservable instantaneous rate, LMM captures the dynamics of the entire curve of interest rates by using the (market observable) LIBOR forwards as its state variables. The time evolution

of the forwards is given by a set of intuitive stochastic differential equations in a way which guarantees arbitrage freeness of the process. The model is intrinsically multi-factor, meaning that it captures accurately various aspects of the curve dynamics: parallel shifts, steepenings / flattenings, butterflies, etc. In this lecture we discuss two versions of the LMM methodology: (i) the classic LMM with a local volatility specification, and (ii) its SABR style extension.

One of the main difficulties experienced by the pre-LMM term structure models is the fact that they tend to produce unrealistic volatility structures of forward rates. The persistent "hump" occurring in the short end of the volatility curve leads to overvaluation of instruments depending on forward volatility. The LMM model offers a solution to this problem by allowing one to impose an approximately stationary volatility and correlation structure of LIBOR forwards. This reflects the view that the volatility structure of interest rates retains its shape over time, without distorting the valuation of instruments sensitive to forward volatility.

On the downside, LMM is far less tractable than, for example, the Hull-White model. In addition, it is not Markovian in the sense short rate models are Markovian. As a consequence, all valuations based on LMM have to be done by means of Monte Carlo simulations.

2   LIBOR market model

2.1   Dynamics of the LIBOR market model

We shall consider a sequence of approximately equally spaced dates $0 = T_0 < T_1 < \ldots < T_N$ which will be termed the *standard tenors*. A standard LIBOR forward rate L_j, $j = 0, 1, \ldots, N - 1$ is associated with a FRA which starts on T_j and matures on T_{j+1}. Usually, it is assumed that $N = 120$ and the L_j's are 3 month LIBOR forward rates. Note that these dates refer to the actual start and end dates of the contracts rather than the LIBOR "fixing dates", i.e. the dates on which the LIBOR rates settle. To simplify the notation, we shall disregard the difference between the contract's start date and the corresponding forward rate's fixing date. Proper implementation, however, must take this distinction into account.

Each LIBOR forward L_j is modeled as a continuous time stochastic process $L_j(t)$. Clearly, this process has the property that it gets killed at $t = T_j$. The dynamics of the forward process is driven by an N-dimensional, correlated Wiener process $W_1(t), \ldots, W_N(t)$. We let ρ_{jk} denote the correlation coefficient between

$W_j(t)$ and $W_k(t)$:

$$\mathsf{E}\left[dW_j(t)dW_k(t)\right] = \rho_{jk}dt\,,$$

where E denotes expected value.

In order to motivate the form of the stochastic differential equations describing the dynamics of the LIBOR forwards, let us first consider the world in which there is no volatility of interest rates. The shape of the forward curve would be set once and for all by a higher authority, and each LIBOR forward would have a constant value $L_j(t) = L_{j0}$. In other words,

$$dL_j(t) = 0,$$

for all j's. The fact that the rates are stochastic forces us to replace this simple dynamical system with a system of stochastic differential equations of the form:

$$dL_j(t) = \Delta_j(t)dt + C_j(t)dW_j(t), \tag{1}$$

where

$$\Delta_j(t) = \Delta_j(t, L(t)),$$
$$C_j(t) = C_j(t, L(t)),$$

are the drift and instantaneous volatility, respectively.

As discussed in Lecture 3, the *no arbitrage* requirement of asset pricing forces a relationship between the drift term and the diffusion term: the form of the drift term depends thus on the choice of numeraire.

Recall from Lecture 3 that L_k is a martingale under the T_{k+1}-forward measure Q_k, and so its dynamics reads:

$$dL_k(t) = C_k(t)dW_k(t),$$

where $C_k(t)$ is an instantaneous volatility function which will be defined later. For $j \neq k$,

$$dL_j(t) = \Delta_j(t)dt + C_j(t)dW_j(t).$$

Since the j-th LIBOR forward settles at T_j, the process for $L_j(t)$ is killed at $t = T_j$. We shall determine the drifts $\Delta_j(t)$ by requiring lack of arbitrage.

Let us first assume that $j < k$. The numeraires for the measures Q_j and Q_k are the prices $P(t, T_{j+1})$ and $P(t, T_{k+1})$ of the zero coupon bonds expiring at T_{j+1} and T_{k+1}, respectively. Explicitly,

$$P(t, T_{j+1}) = P(t, T_{\gamma(t)}) \prod_{\gamma(t) \leq i \leq j} \frac{1}{1 + \delta_i F_i(t)}\,, \tag{2}$$

where F_i denotes the OIS forward[1] spanning the accrual period $[T_i, T_{i+1})$, and where $\gamma : [0,\ T_N] \to \mathbb{Z}$ is defined by

$$\gamma(t) = m + 1, \quad \text{if } t \in [T_m,\ T_{m+1})\ .$$

Notice that $P(t, T_{\gamma(t)})$ is the "stub" discount factor over the incomplete accrual period $[t, T_{\gamma(t)}]$.

Since the drift of $L_j(t)$ under Q_j is zero, formula (26) (or (27)) of Lecture 2 yields:

$$\begin{aligned}
\Delta_j(t) &= \frac{d}{dt}\left[L_j,\ \log \frac{P(\,\cdot\,, T_{j+1})}{P(\,\cdot\,, T_{k+1})} \right](t) \\
&= -\frac{d}{dt}\left[L_j,\ \log \prod_{j+1 \leq i \leq k} (1 + \delta_i F_i) \right](t) \\
&= -\sum_{j+1 \leq i \leq k} dL_j(t)\, \frac{\delta_i dF_i(t)}{1 + \delta_i F_i(t)} \\
&= -C_j(t) \sum_{j+1 \leq i \leq k} \frac{\rho_{ji}\delta_i C_i(t)}{1 + \delta_i F_i(t)}\ ,
\end{aligned}$$

where, in the third line, we have used the fact that the spread between L_j and F_j is deterministic, and thus its contribution to the quadratic variation is zero.

Similarly, for $j > k$, we find that

$$\Delta_j(t) = C_j(t) \sum_{k+1 \leq i \leq j} \frac{\rho_{ji}\delta_i C_i(t)}{1 + \delta_i F_i(t)}\ .$$

We can thus summarize the above discussion as follows. In order to streamline the notation, we let $dW_j(t) = dW_j^{\mathsf{Q}_k}(t)$ denote the Wiener process under the measure Q_k. Then the dynamics of the LMM model is given by the following system of stochastic differential equations. For $t < \min(T_k, T_j)$,

$$dL_j(t) = C_j(t) \\
\times \begin{cases}
-\sum_{j+1 \leq i \leq k} \dfrac{\rho_{ji}\delta_i C_i(t)}{1 + \delta_i F_i(t)}\, dt + dW_j(t), & \text{if } j < k, \\[2ex]
dW_j(t), & \text{if } j = k, \\[2ex]
\sum_{k+1 \leq i \leq j} \dfrac{\rho_{ji}\delta_i C_i(t)}{1 + \delta_i F_i(t)}\, dt + dW_j(t), & \text{if } j > k\ .
\end{cases} \tag{3}$$

---

[1] Recall that all discounting is done on OIS.

These equations have to be supplied with initial values for the LIBOR forwards:

$$L_j(0) = L_{j0}, \tag{4}$$

where L_{j0} is the current value of the forward which is implied by the current yield curve.

In addition to the forward measures discussed above, it is convenient to use the spot measure. It is expressed in terms of the numeraire:

$$B(t) = \frac{P(t, T_{\gamma(t)})}{\prod_{1 \leq i \leq \gamma(t)} P(T_{i-1}, T_i)} . \tag{5}$$

Under the spot measure, the LMM dynamics reads:

$$dL_j(t) = C_j(t) \Big(\sum_{\gamma(t) \leq i \leq j} \frac{\rho_{ji} \delta_i C_i(t)}{1 + \delta_i F_i(t)} \, dt + dW_j(t) \Big). \tag{6}$$

2.2 Structure of the instantaneous volatility

So far we have been working with a general instantaneous volatility $C_j(t)$ for the forward $L_j(t)$. In practise, we assume $C_j(t)$ to be one of the following standard volatility specifications discussed in Lecture 3:

$$C_j(t) = \begin{cases} \sigma_j(t) & \text{(normal model)}, \\ \sigma_j(t) L_j(t)^{\beta_j} & \text{(CEV model)}, \\ \sigma_j(t) L_j(t) & \text{(lognormal model)}, \\ \sigma_j(t) L_j(t) + \vartheta_j(t) & \text{(shifted lognormal model)}, \end{cases} \tag{7}$$

where the functions $\sigma_j(t)$ and $\vartheta_j(t)$ are deterministic, and where $\beta_j \leq 1$. In the following, we will be assuming the CEV model specification, and thus the dynamics of the LIBOR forwards is given by the system:

$$dL_j(t) = \sigma_j(t) L_j(t)^{\beta_j}$$
$$\times \begin{cases} -\sum_{j+1 \leq i \leq k} \dfrac{\rho_{ji} \delta_i \sigma_i(t) L_i(t)^{\beta_i}}{1 + \delta_i F_i(t)} \, dt + dW_j(t), & \text{if } j < k, \\ dW_j(t), & \text{if } j = k, \\ \sum_{k+1 \leq i \leq j} \dfrac{\rho_{ji} \delta_i \sigma_i(t) L_i(t)^{\beta_i}}{1 + \delta_i F_i(t)} \, dt + dW_j(t), & \text{if } j > k, \end{cases} \tag{8}$$

under Q_k, or under the spot measure:

$$dL_j(t) = \sigma_j(t)L_j(t)^{\beta_j} \Big(\sum_{\gamma(t) \le i \le j} \frac{\rho_{ji}\delta_i\sigma_i(t)L_i(t)^{\beta_i}}{1 + \delta_i F_i(t)} \, dt + dW_j(t) \Big). \qquad (9)$$

We recall from the discussion in Lecture 3 that the CEV model needs care at zero forward. Experience shows that the Dirichlet (absorbing) boundary condition at zero works better than the Neumann (reflecting) condition, and we will assume that the Dirichlet condition is imposed. What it means is that if a path realizing the process for L_j hits zero, it gets killed and stays zero forever.

2.3   Factor reduction

In a market where the forward curve spans 30 years, there are 120 quarterly LIBOR forwards and thus 120 stochastic factors. So far we have not imposed any restrictions on the number of these factors, and thus the number of Brownian motions driving the LIBOR forward dynamics is equal to the number of forwards. Having a large number of factors poses severe problems with the model's implementation. On the numerical side, the "curse of dimensionality" kicks in, leading to unacceptably slow performance. On the financial side, the parameters of the model are severely underdetermined and the calibration of the model becomes unstable.

   We are thus led to the assumption that only a small number d of independent Brownian motions $Z_a(t)$, $a = 1, \ldots, d$, with

$$\mathsf{E}\left[dZ_a(t)dZ_b(t)\right] = \delta_{ab}dt \,, \qquad (10)$$

should drive the process. Typically, $d = 1, 2, 3$, or 4. We set

$$dW_j(t) = \sum_{1 \le a \le d} U_{ja} \, dZ_a(t), \qquad (11)$$

where U is an $N \times d$ matrix with the property that UU' is close to the correlation matrix. Of course, it is in general impossible to have $UU' = \rho$. We can easily rewrite the dynamics of the model in terms of the independent Brownian motions:

$$dL_j(t) = \Delta_j(t)dt + \sum_{1 \le a \le d} B_{ja}(t)dZ_a(t), \qquad (12)$$

where

$$B_{ja}(t) = U_{ja} \, C_j(t). \qquad (13)$$

We shall call this system the *factor reduced* LMM dynamics. It is the factor reduced form of LMM that is used in practice.

3   Calibration of the LMM model

Calibration (to a selected collection of benchmark instruments) is a choice of the model parameters so that the model reprices the benchmark instruments to a desired accuracy. The choice of the calibrating instruments is dictated by the characteristics of the portfolio to be managed by the model.

   An important feature of LMM is that it leads to pricing formulas for caps and floors which are consistent with the market practice of quoting the prices of these products in terms of Black's model. This makes the calibration of LMM to caps and floors very easy. On the other hand, from the point of view of the LMM model, swaptions are exotic structures whose fast pricing poses serious challenges. In this section we describe our strategy of dealing with these issues.

3.1   Approximate valuation of swaptions

A key ingredient of any efficient calibration methodology for LMM is rapid and accurate swaption valuation. A swap rate is a non-linear function of the underlying LIBOR forward rates. The stochastic differential equation for the swap rate implied by the LMM model cannot be solved in closed form, and thus pricing swaptions within LMM requires Monte Carlo simulations. This poses a serious issue for efficient model calibration, as such simulations are very time consuming.

   Let us describe a closed form approximation which can be used to calibrate the model. We consider a standard forward starting swap, whose start and end dates are denoted by T_m and T_n, respectively. Recall from Lecture 1 that the level function of the swap is defined by:

$$A_{mn}(t) = \sum_{m \leq j \leq n-1} \alpha_j P\left(t, T_{j+1}\right), \tag{14}$$

where α_j are the day count fractions for fixed rate payments, and where $P\left(t, T_j\right)$ is the time t value of \$1 paid at time T_j. Typically, the payment frequency on the fixed leg is not the same as that on the floating leg[2] (which we continue to denote

---

[2]Remember, the default convention on US dollar swaps is a semiannual $30/360$ fixed leg versus a quarterly $\mathrm{act}/360$ floating leg.

by δ_j). This fact causes a bit of a notational nuisance but needs to be taken properly into account for accurate pricing. We let $S_{mn}(t)$ denote the corresponding forward swap rate. In order to lighten up the notation, we will suppress the subscripts mn throughout the remainder of this lecture.

A straightforward calculation shows that, under the forward measure Q_k, the dynamics of the swap rate process can be written in the form:

$$dS(t) = \Omega(t, L)dt + \sum_{m \le j \le n-1} \Lambda_j(t, L)dW_j(t), \tag{15}$$

where

$$\Omega = \sum_{m \le j \le n-1} \frac{\partial S}{\partial F_j} \Delta_j + \frac{1}{2} \sum_{m \le i,j, \le n-1} \rho_{ij} \frac{\partial^2 S}{\partial F_i \partial F_j} C_i C_j, \tag{16}$$

and

$$\Lambda_j = \frac{\partial S}{\partial F_j} C_j. \tag{17}$$

Not surprisingly, the stochastic differential equation for S has a drift term: the forward swap rate is not a martingale under a forward measure. Shifting to the martingale measure Q_{mn} (the swap measure),

$$dW(t) = \frac{\sum_{m \le j \le n-1} \Lambda_j(t, F)dW_j(t) + \Omega(t, F)dt}{\nu_{mn}(t)}, \tag{18}$$

where

$$\nu_{mn}(t)^2 = \sum_{m \le i,j \le n-1} \rho_{ij}\Lambda_i(t, L)\Lambda_j(t, L), \tag{19}$$

we get

$$dS(t) = \nu(t)dW(t). \tag{20}$$

In order to be able to use this dynamics effectively, we have to approximate it by quantities with tractable analytic forms. The simplest approximation consists in replacing the values of the stochastic forwards $L_j(t)$ by their initial values L_{j0}. This amounts to "freezing" the curve at its current shape. Within this approximation, the coefficients in the diffusion process (15) for the swap rate are deterministic:

$$\Lambda_j(t, L) \approx \Lambda_j(t, L_0), \tag{21}$$

and

$$\Omega(t, L) \approx \Omega(t, L_0). \tag{22}$$

Let $\nu_0(t)$ denote the value of $\nu(t)$ in this approximation, i.e. $\nu_0(t)$ is given by (19) with all $\Lambda_j(t, L)$ replaced by $\Lambda_j(t, L_0)$. The stochastic differential equation (20) can then be solved in closed form,

$$S(t) = S_0 + \int_0^t \nu_0(s)dW(s). \tag{23}$$

This is a normal model with deterministic time dependent volatility and thus the swaption implied normal volatility ζ_{mn} is approximately given by

$$
\begin{aligned}
\zeta_{mn}^2 &\approx \frac{1}{T_m} \int_0^t \nu_0(s)^2 ds \\
&= \frac{1}{T_m} \sum_{m \le j, l \le n-1} \rho_{jl} \int_0^t \Lambda_j(s, F_0)\Lambda_l(s, F_0)ds \ .
\end{aligned}
\tag{24}
$$

This formula is easy to implement in code, and leads to reasonably accurate results.

The frozen curve approximation can be regarded as the lowest order term in the "small noise expansion". With a bit of extra work, one can compute higher order terms in that expansion.

3.2   Parametrization of the volatility surface

For the purpose of calibration we require that the deterministic instantaneous CEV volatilities $\sigma_j(t)$ in (8) are piecewise constant. In order to help the intuition, we organize constant components as a lower triangular matrix in Table 1.

	$t \in [T_0, T_1)$	$t \in [T_1, T_2)$	$\dots$	$t \in [T_{N-1}, T_N)$
$\sigma_0(t)$	0	0	$\dots$	0
$\sigma_1(t)$	$\sigma_{1,0}$	0	$\dots$	0
$\sigma_2(t)$	$\sigma_{2,0}$	$\sigma_{2,1}$	$\dots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\sigma_{N-1}(t)$	$\sigma_{N-1,0}$	$\sigma_{N-1,1}$	$\dots$	0

Table 1: General volatility structure

Clearly, the problem of determining all the $\sigma_{j,i}$'s is vastly overparametrized. Table 1 contains 7140 parameters (assuming $N = 120$)! A natural remedy to

the overparametrization problem is assuming that the instantaneous volatility is *stationary*, i.e.,

$$\begin{aligned}
\sigma_{j,i} &= \sigma_{j-i,0} \\
&\equiv \sigma_{j-i},
\end{aligned} \tag{25}$$

for all $i < j$. This assumption appears natural and intuitive, as it implies that the structure of cap volatility will look in the future exactly the same way as it does currently. Consequently, the "forward volatility problem" plaguing the traditional terms structure models would disappear. Under the stationary volatility assumption, the instantaneous volatility has the structure summarized in Table 2. It is a good idea to reduce the number of parameters even further, and try to find

	$t \in [T_0, \, T_1)$	$t \in [T_1, \, T_2)$	$\ldots$	$t \in [T_{N-1}, \, T_N)$
$\sigma_0(t)$	0	0	$\ldots$	0
$\sigma_1(t)$	σ_1	0	$\ldots$	0
$\sigma_2(t)$	σ_2	σ_1	$\ldots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\sigma_{N-1}(t)$	σ_{N-1}	σ_{N-2}	$\ldots$	0

Table 2: Stationary volatility structure

a parametric fit $\sigma_i = h(T_i)$, $i = 1, \ldots, N - 1$. A popular (but, by no means the only) choice is the hump function

$$h(t) = (at + b)e^{-\lambda t} + \mu. \tag{26}$$

Despite its intuitive appeal, the stationarity assumption is not sufficient for accurate calibration of the model. The financial reason behind this fact appears to be the phenomenon of mean reversion of long term rates. Unlike the Vasicek style models, it is impossible to take this phenomenon into account by adding an Ornstein - Uhlenbeck style drift term to the LMM dynamics as this would violate the arbitrage freeness of the model. On the other hand, one can achieve a similar effect by suitably specifying the instantaneous volatility function. In order to implement this idea, we assume that the long term volatility structure is given by $\overline{\sigma}_i = \overline{h}(T_i)$, $i = 1, \ldots, N - 1$, where $\overline{h}(t)$ is another hump shaped function. We then set

$$\sigma_{j,i} = p_i \sigma_{j-i} + q_i \overline{\sigma}_{j-i}, \tag{27}$$

i.e. the σ's are mixtures of the short term σ's and the equilibrium $\overline{\sigma}$'s. The weights p_i and q_i are parametrized so that $p_i, q_i \geq 0$, $p_i + q_i = 1$, and $p_i \to 0$, as $i \to \infty$. In other words, as we move forward in time, the volatility structure looks more and more like the long term limit. This specification is summarized in Table 3.

	$t \in [T_0,\, T_1)$	$t \in [T_1,\, T_2)$	$\ldots$	$t \in [T_{N-1},\, T_N)$
$\sigma_0(t)$	0	0	$\ldots$	0
$\sigma_1(t)$	$p_1\sigma_1 + q_1\overline{\sigma}_1$	0	$\ldots$	0
$\sigma_2(t)$	$p_1\sigma_2 + q_1\overline{\sigma}_2$	$p_2\sigma_1 + q_2\overline{\sigma}_1$	$\ldots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\sigma_{N-1}(t)$	$p_1\sigma_{N-1} + q_1\overline{\sigma}_{N-1}$	$p_2\sigma_{N-2} + q_2\overline{\sigma}_{N-2}$	$\ldots$	0

Table 3: Approximately stationary volatility structure

The lower triangular matrix in Table 3, LMM's internal representation of volatility, is referred to as the *LMM volatility surface*. We leave out the details of this methodology, as that would make the presentation a bit tedious. In the final result, we have a parametrization of the volatility surface by a manageable number of parameters $\theta = (\theta_1, \ldots, \theta_d)$ (such as the parameters of the hump functions $h(t)$ and $\overline{h}(t)$, and of the weights p_i), such that $\sigma_{j,i} = \sigma_{j,i}(\theta)$ can be calibrated to the market and has an intuitive shape.

3.3   Parametrization of the correlation matrix

The central issue is to calibrate the model, at the same time, to the cap / floor and swaption markets in a stable and consistent manner. An important part of this process is determining the correlation matrix $\rho = \{\rho_{jk}\}_{0 \leq j,k \leq N-1}$. The dimensionality of ρ is $N(N+1)/2$, clearly far too high to assure a stable calibration procedure.

A convenient approach to correlation modeling is to use a parameterized form of ρ_{ij}. An intuitive and flexible parametrization is given by the formula:

$$\rho_{ij} = \overline{\rho}_{\min(i,j)} + \left(1 - \overline{\rho}_{\min(i,j)}\right) \exp\left(-\beta_{\min(i,j)} |T_i - T_j|\right), \tag{28}$$

where

$$\overline{\rho}_k = \rho \tanh\left(\alpha T_k\right), \tag{29}$$

and

$$\beta_k = \beta T_k^{-\kappa}. \tag{30}$$

The meaning of the parameters is as follows: ρ is the asymptotic level of correlations, α is a measure of speed at which ρ is approached, β is a the decay rate of correlations, and κ is an asymmetry parameter. Intuitively, positive κ means that two consecutive forwards with short maturities are less correlated than two such forwards with long maturities. The parameters in this formula can be calibrated by using, for example, historical data. A word of caution is in order: this parametrization produces a matrix that is only approximately positive definite.

3.4   Optimization

In order to calibrate the model we seek instantaneous volatility parameters σ_i so that to fit the at the money caplet and swaption volatilities. These can be expressed in terms of the instantaneous volatilities are as follows.

Let ζ_m denote the at the money implied normal volatility of the caplet expiring at T_m. Then, within the frozen curve approximation,

$$\zeta_m(\theta)^2 = \frac{1}{T_m} \, L_{m0}^{2\beta_m} \sum_{0 \le i \le m-1} \sigma_{m,i}(\theta)^2 \, \delta_i, \tag{31}$$

where $\delta_i = T_{i+1} - T_i$, and θ denotes the set of parameters of the LMM volatility surface. This relationship is reasonably accurate, and can be used for calibration. However, in practice, one needs to improve on this formula by going beyond the frozen curve approximation.

Similarly, for the at the money implied normal volatility ζ_{mn} of the swaption expiring at T_m into a swap maturing at T_n we have an approximate expression:

$$\zeta_{mn}(\theta)^2 = \frac{1}{T_m} \sum_{0 \le i \le m-1} \sum_{m \le j,l \le n-1} \rho_{jl} \Lambda_{j,i} \Lambda_{l,i} \delta_i, \tag{32}$$

where $\Lambda_{j,i}$ is the (constant) value of $\Lambda_j(s, L_0)$, the frozen curve approximation to (17), for $s \in [T_i, T_{i+1})$. Note that the coefficients $\Lambda_{j,i}$ depend on the parameters of the LMM volatility surface.

The objective function for optimization is given by:

$$\mathcal{L}(\theta) = \sum_m w_m \left(\zeta_m(\theta) - \overline{\zeta}_m \right)^2 + \sum_{m,n} w_{mn} \left(\zeta_{mn}(\theta) - \overline{\zeta}_{mn} \right)^2, \tag{33}$$

where $\overline{\zeta}_m$ and $\overline{\zeta}_{mn}$ are the market observed caplet and swaption implied normal volatilities. The coefficients w_m and w_{mn} are weights which allow the user select

the calibration instruments and their relative importance. Finally, it is a good idea to add a Tikhonov style regularization in order to maintain stability of the calibration. A convenient and computationally efficient choice of the Tikhonov penalty term is the integral of the square of the mean curvature (of elementary differential geometry of surfaces) of the parameterized LMM volatility surface. The impact of this penalty term is to discourage regions of extreme curvature (such as a sharp ridge along the diagonal) at the expense of slightly less accurate fit.

4   Generating Monte Carlo paths for LMM

LMM does not allow for a natural implementation based on recombining trees, and thus all valuations have to be performed via Monte Carlo simulations. We shall describe two numerical schemes for generating Monte Carlo paths for LMM: Euler's scheme and Milstein's scheme. They both consist in replacing the infinitesimal differentials by suitable finite differences.

We choose a sequence of event dates $t_0, t_1, \ldots, t_m$, and denote by $L_{jn} \simeq L_j(t_n)$ the approximate solution. We also set

$$
\begin{aligned}
\Delta_{jn} &= \Delta_j(t_n, L_n), \\
B_{jan} &= B_{ja}(t_n, L_{jn}),
\end{aligned}
\tag{34}
$$

and $\delta t_n = t_{n+1} - t_n$.

Euler's scheme. This is the simplest, universally applicable discretization method. Applied to LMM, it reads:

$$
L_{j,n+1} = L_{jn} + \Delta_{jn}\delta t_n + \sum_{1 \leq a \leq d} B_{jan}\,\delta Z_{na} \, ,
\tag{35}
$$

where $\delta Z_{na} = Z_a(t_{n+1}) - Z_a(t_n)$ is the discretized Brownian motion, see Appendix. Euler's scheme is of order of convergence $1/2$ meaning that the approximate solution converges in a suitable norm to the actual solution at the rate of $\delta t^{1/2}$, as $\delta t \equiv \max \delta t_n \to 0$.

Milstein's scheme. This is a refinement of Euler's scheme. It is universally applicable to stochastic differential equations driven by a single Brownian motion, but generally does not work for equations driven by several Brownian motions (for example, it is not suitable for the SABR model). Fortunately, LMM is in the category of models which satisfy the assumptions required for Milstein's scheme to work.

In order to lighten up the notation, let us define:

$$\Upsilon_{jabn} \equiv B_{ja}(t_n, L_{jn}) \frac{\partial B_{jb}(t_n, L_{jn})}{\partial L_j}.$$ (36)

Then Milstein's scheme for the LMM model reads:

$$
\begin{aligned}
L_{j,n+1} = L_{jn} &+ \left(\Delta_{jn} - \frac{1}{2} \sum_{1 \leq a \leq d} \Upsilon_{jaan}\right)\delta t_n \\
&+ \sum_{1 \leq a \leq d} B_{jan}\,\delta Z_{na} + \frac{1}{2} \sum_{1 \leq a,b \leq d} \Upsilon_{jabn}\,\delta Z_{na}\,\delta Z_{nb}.
\end{aligned}
$$ (37)

Milstein's scheme is of order of convergence 1 meaning that the approximate solution converges in a suitable norm to the actual solution at the rate of δt, as $\delta t \to 0$.

A bit of a challenge lies in handling the drift terms. Because of their complexity, their calculation (at each time step) takes up to 50% of the total computation time. A simple remedy to this problem is to freeze the drift terms at today's values of the forward curve (the frozen curve approximation) but this leads to noticeable inaccuracies in pricing of longer dated options. Going one step in the low noise expansion beyond the frozen curve approximation produces satisfying results.

5 The SABR / LMM model

The classic LMM model has a severe drawback: while it is possible to calibrate it so that it matches at the money option prices, it generally misprices out of the money options. The main reason for this is its specification. While the market uses stochastic volatility models in order to price out of the money vanilla options, LMM is incompatible with such models. In order to remedy the problem, we describe a model that combines the key features of the LMM and SABR models.

5.1 Dynamics of the SABR / LMM model

To this end, we assume that the instantaneous volatilities $C_j(t)$ of the forward rates L_j are of the form

$$C_j(t) = \sigma_j(t)L_j(t)^{\beta_j},$$ (38)

with stochastic volatility parameters $\sigma_j(t)$. Furthermore, we assume that, under the T_{k+1}-forward measure $\mathbb{Q}_k$, the full dynamics of the forward is given by the

stochastic system:

$$dL_k(t) = C_k(t)dW_k(t),$$
$$d\sigma_k(t) = D_k(t)dZ_k(t),$$
(39)

where the diffusion coefficient of the process $\sigma_k(t)$ is of the form

$$D_k(t) = \alpha_k(t)\sigma_k(t).$$
(40)

Note that $\alpha_k(t)$ is assumed here to be a (deterministic) function of t rather than a constant. This extra flexibility is added in order to make sure that the model can be calibrated to market data.

In addition, we impose the following instantaneous volatility structure:

$$\mathsf{E}[dW_j(t)dZ_k(t)] = r_{jk}\,dt,$$
(41)

and

$$\mathsf{E}[dZ_j(t)dZ_k(t)] = \eta_{jk}dt.$$
(42)

The block matrix

$$\Pi = \begin{bmatrix} \rho & r \\ r' & \eta \end{bmatrix}$$
(43)

is assumed to be positive definite.

Let us now derive the dynamics of such an extended LIBOR market model under the common forward measure Q_k. According to the arbitrage pricing theory, the form of the stochastic differential equations defining the dynamics of the LIBOR forward rates depends on the choice of numeraire.

Under the T_{k+1}-forward measure Q_k, the dynamics of the forward rate $L_j(t)$, $j \neq k$ reads:

$$dL_j(t) = \Delta_j(t)dt + C_j(t)dW_j(t).$$

We determine the drifts $\Delta_j(t) = \Delta_j(t, L(t), \sigma(t))$ by requiring lack of arbitrage. This is essentially the same calculation as in the derivation of the drift terms for the classic LMM, and we can thus summarize the result as follows. In order to streamline the notation, we let $dW(t) = dW^{\mathsf{Q}_k}(t)$ denote the Wiener process under the measure Q_k. Then, as expected,

$$dL_j(t) = C_j(t)$$
$$\times \begin{cases} -\sum_{j+1 \leq i \leq k} \dfrac{\rho_{ji}\delta_i C_i(t)}{1 + \delta_i F_i(t)}\,dt + dW_j(t), & \text{if } j < k, \\ dW_j(t), & \text{if } j = k, \\ \sum_{k+1 \leq i \leq j} \dfrac{\rho_{ji}\delta_i C_i(t)}{1 + \delta_i F_i(t)}\,dt + dW_j(t), & \text{if } j > k\,. \end{cases}$$
(44)

Similarly, under the spot measure, the SABR / LMM dynamics reads:

$$dL_j(t) = C_j(t)\Big(\sum_{\gamma(t)\leq i\leq j} \frac{\rho_{ji}\delta_i C_i(t)}{1+\delta_i F_i(t)}\, dt + dW_j(t) \Big). \qquad (45)$$

Let us now compute the drift term $\Gamma_j(t) = \Gamma_j(t, L(t), \sigma(t))$ for the dynamics of $\sigma_j(t)$, $j \neq k$, under Q_k,

$$d\sigma_j(t) = \Gamma_j(t)dt + D_j(t)dZ_j(t).$$

Let us first assume that $j < k$. The numeraires for the measures Q_j and Q_k are the prices $P(t, T_{j+1})$ and $P(t, T_{k+1})$ of the zero coupon bonds maturing at T_{j+1} and T_{k+1}, respectively. Since the drift of $L_j(t)$ under Q_j is zero, formula (26) of Lecture 2 yields:

$$\begin{aligned}
\Gamma_j(t) &= \frac{d}{dt}\Big[\sigma_j,\ \log\frac{P(\cdot, T_{j+1})}{P(\cdot, T_{k+1})}\Big](t)\\
&= -\frac{d}{dt}\Big[\sigma_j,\ \log\prod_{j+1\leq i\leq k}(1+\delta_i F_i)\Big](t)\\
&= -\sum_{j+1\leq i\leq k} d\sigma_j(t)\,\frac{\delta_i dF_i(t)}{1+\delta_i F_i(t)}\\
&= -D_j(t)\sum_{j+1\leq i\leq k}\frac{r_{ji}\delta_i C_i(t)}{1+\delta_i F_i(t)}\, dt.
\end{aligned}$$

Similarly, for $j > k$, we find that

$$\Gamma_j(t) = D_j(t)\sum_{k+1\leq i\leq j}\frac{r_{ji}\delta_i C_i(t)}{1+\delta_i F_i(t)}\ .$$

This leads to the following system:

$$d\sigma_j(t) = D_j(t)$$

$$\times \begin{cases}
-\sum_{j+1\leq i\leq k}\dfrac{r_{ji}\delta_i C_i(t)}{1+\delta_i F_i(t)}\, dt + dZ_j(t), & \text{if } j < k,\\[2mm]
dZ_j(t), & \text{if } j = k,\\[2mm]
\sum_{k+1\leq i\leq j}\dfrac{r_{ji}\delta_i C_i(t)}{1+\delta_i F_i(t)}\, dt + dZ_j(t), & \text{if } j > k,
\end{cases} \qquad (46)$$

under Q_k, and

$$d\sigma_j(t) = D_j(t)\Big(\sum_{\gamma(t)\le i\le j} \frac{r_{ji}\delta_i D_i(t)}{1+\delta_i F_i(t)}\, dt + dZ_j(t) \Big), \tag{47}$$

under the spot measure.

We now plug in the explicit choices made in (38) and (40). Under the T_{k+1}-forward measure Q_k, the dynamics of the full model reads:

$$dF_j(t) = \sigma_j(t)L_j(t)^{\beta_j}$$
$$\times \begin{cases} -\sum_{j+1\le i\le k} \dfrac{\rho_{ji}\delta_i\sigma_i(t)L_i(t)^{\beta_i}}{1+\delta_i F_i(t)}\, dt + dW_j(t), & \text{if } j < k, \\[2ex] dW_j(t), & \text{if } j = k, \\[2ex] \sum_{k+1\le i\le j} \dfrac{\rho_{ji}\delta_i\sigma_i(t)L_i(t)^{\beta_i}}{1+\delta_i F_i(t)}\, dt + dW_j(t), & \text{if } j > k\,. \end{cases} \tag{48}$$

and

$$d\sigma_j(t) = \alpha_j(t)\sigma_j(t)$$
$$\times \begin{cases} -\sum_{j+1\le i\le k} \dfrac{r_{ji}\delta_i\sigma_i(t)L_i(t)^{\beta_i}}{1+\delta_i F_i(t)}\, dt + dZ_j(t), & \text{if } j < k, \\[2ex] dZ_j(t), & \text{if } j = k, \\[2ex] \sum_{k+1\le i\le j} \dfrac{r_{ji}\delta_i\sigma_i(t)L_i(t)^{\beta_i}}{1+\delta_i F_i(t)}\, dt + dZ_j(t), & \text{if } j > k\,, \end{cases} \tag{49}$$

supplemented by the initial conditions:

$$\begin{aligned} L_j(0) &= L_{j0}, \\ \sigma_j(0) &= \sigma_{j0}. \end{aligned} \tag{50}$$

Here, L_{j0}'s and σ_{j0}'s are the currently observed values. Similarly, under the spot measure Q_0, the dynamics is given by the stochastic system:

$$dL_j(t) = \sigma_j(t)L_j(t)^{\beta_j}\Big(\sum_{\gamma(t)\le i\le j} \frac{\rho_{ji}\delta_i\sigma_i(t)L_i(t)^{\beta_i}}{1+\delta_i F_i(t)}\, dt + dW_j(t) \Big),$$
$$d\sigma_j(t) = \alpha_j(t)\sigma_j(t)\Big(\sum_{\gamma(t)\le i\le j} \frac{r_{ji}\delta_i\sigma_i(t)L_i(t)^{\beta_i}}{1+\delta_i F_i(t)}\, dt + dZ_j(t) \Big). \tag{51}$$

5.2 Practicalities of the SABR / LMM model

Time does not permit us to get into any detailed discussion of the practical aspects of SABR / LMM, and we will just highlight a number of issues. To a large degree, in order to implement SABR / LMM one follows the steps described above in the case of the classic LMM. The model has to be factor reduced in order to make it practical, sensible parametrizations for the volatilities and correlations have to be found, and a good deal of analytic work needs to be done to prepare ground for calibration. Let us note a number of new features of the SABR / LMM model as compare to the original models.

Not surprisingly, unlike the classic LMM model, exact closed form valuation of caps and floors is not possible in SABR / LMM. This is simply a reflection of the fact that SABR itself does not have have closed form solutions, and one either relies on sensible approximations or Monte Carlo simulations. However, the reassuring fact is that SABR / LMM allows for pricing of caps / floors which is in principle consistent with market practice. This can be seen as follows. Assume that we have chosen the T_{k+1}-forward measure Q_k for pricing. A cap is a basket of caplets spanning a number of consecutive accrual periods. Consider the caplet spanning the period $[T_j, T_{j+1}]$. Shifting from Q_k to the T_{j+1}-forward measure Q_j, we note that its dynamics is that of the classic SABR model. Since instrument valuation is invariant under change of numeraire, this shows that the price of the caplet is consistent with its SABR price.

The correlation structure of SABR / LMM is very rich: in addition to the block of correlations between the forwards, we have the blocks of correlations between the volatilities, and the block of correlations between the forwards and volatilites. Together, these correlations determine the shape of volatility smile.

SABR / LMM specifies the values of the CEV exponents β_j for each benchmark forward L_j but it does not use explicit CEV exponents β_{mn} for the benchmark forward swap rates S_{mn}. These are internally implied by the model. There is no simple relation between the caplet β's and the swaption β's. An approximation which works well in practice is given by the following formulas:

$$\beta_{mn} = \sum_{m \leq j \leq n-1} a_{mn,j}\beta_j + b_{mn}, \qquad (52)$$

where

$$a_{mn,k} = \frac{2 \log L_{k0}}{(n-m)^2} \sum_{m \le j \le n-1} \frac{1}{\log L_{j0} + \log L_{k0}} \ ,$$

$$b_{mn} = \frac{1}{(n-m)^2} \sum_{m \le j,k \le n-1} \frac{\log \rho_{jk}}{\log L_{j0} + \log L_{k0}} \ . \tag{53}$$

Note that

$$\sum_{m \le j \le n-1} a_{mn,j} = 1. \tag{54}$$

Consequently, the CEV power of a swaption is a weighted average of the CEV powers of the spanning forwards plus a convexity correction. Under a perfectly flat forward curve $a_{mn,j} = 1/(n-m)$, for all j. The convexity correction b is rather small. On a typical market snapshot it is of the order of magnitude 10^{-3}, and thus for all practical purposes it can be assumed zero.

Swaptions are the most liquid volatility instruments in the interest rates markets, and a term structure model should be calibrated to a suitable set of swaptions. Calibration of SABR / LMM to swaptions requires understanding the relationships between swaption SABR parameters (as discussed in Lecture 3), and the caplet parameters appearing in the SABR / LMM model. An example of such a relationship is the approximate equality (52). Other relations of this type are: relations between caplet β_j-volatility processes $\sigma_j(t)$ and the swaption β_{mn}-volatility processes, relations between the corresponding "volvols", and between the swaption SABR correlation coefficient and the correlation structure of SABR / LMM. Such relationships are fairly easy to derive within the crude "frozen curve" approximation discussed above but, even then, they take some space to write down, and a good deal of coding effort to make them work.

A   Simulating Brownian motion

There exist many more of less refined methods for simulating a Wiener process; here we describe two of them.

The *random walk method* is easy to implement at the expense of being rather noisy. It represents a Wiener process as a random walk sampled at a finite set of event dates $t_0 < t_1 < \ldots < t_m$:

$$W(t_{-1}) = 0,$$

$$W(t_n) = W(t_{n-1}) + \sqrt{t_n - t_{n-1}}\ \xi_n, \quad n = 0, \ldots, m, \tag{55}$$

where $t_{-1} = 0$, and where ξ_n are i.i.d. random variables with $\xi_n \sim N(0, 1)$. A good method of generating the ξ_n's is to first generate a sequence of uniform pseudorandom numbers u_n (using, say, the Mersenne twister algorithm), and then set

$$\xi_n = N^{-1}(u_n), \tag{56}$$

where $N^{-1}(x)$ is the inverse cumulative normal function.

The *spectral decomposition method* generally leads to much better performance than the random walk method. It assures that the simulated process has the same covariance matrix as the Wiener process $W(t)$ sampled at $t_0, t_1, \ldots, t_m$. The latter is explicitly given by:

$$\begin{aligned} \mathcal{C}_{ij} &= \mathsf{E}\big[W(t_i)W(t_j)\big] \\ &= \min(t_i, \, t_j). \end{aligned} \tag{57}$$

Consider the eigenvalue problem for $\mathcal{C}$:

$$\mathcal{C}E_j = \lambda_j E_j, \qquad j = 0, \, \ldots, \, m, \tag{58}$$

with orthonormal E_j's. Since the covariance matrix $\mathcal{C}$ is positive definite, all of its eigenvalues λ_j are nonnegative, and we will assume that

$$\lambda_0 \geq \, \ldots \, \geq \lambda_m \geq 0. \tag{59}$$

We will denote the n-th component of the vector E_j by $E_j(t_n)$, and consider the random variable

$$W(t_n) = \sum_{0 \leq j \leq m} \sqrt{\lambda_j} \, E_j(t_n) \, \xi_j, \tag{60}$$

where ξ_j are, again, i.i.d. random variables with $\xi_j \sim N(0, 1)$. These numbers are best calculated by applying the inverse cumulative normal function to a sequence of Sobol numbers. Alternatively, one could use a sequence of uniform pseudorandom numbers; this, however, leads to a significantly higher sampling variance. Then, for each $n = 0, \, \ldots, \, m$, $W(t_n) \sim N(0, t_n)$, and

$$\begin{aligned} \mathsf{E}\big[W(t_i)W(t_j)\big] &= \sum_{0 \leq k \leq 0} \lambda_k \, E_k(t_i) E_k(t_j) \\ &= \mathcal{C}_{ij}. \end{aligned} \tag{61}$$

We can thus regard $W(t_n)$ a realization of the discretized Wiener process[3]. In practice, we may want to use only a certain portion of the spectral representation

---

[3]This realization of the discretized Wiener process is related to the well known Karhounen-Loeve expansion of the (continuous time) Wiener process.

(60) by truncating it at some $p < m$. This eliminates the *high frequencies* from $W(t_n)$, and lowers the sampling variance. The price for this may be systematically lower accuracy.

References

[1] Brigo, D., and Mercurio, F.: *Interest Rate Models - Theory and Practice*, Springer Verlag (2006).

[2] Hagan, P., and Lesniewski, A.: LIBOR market model with SABR style stochastic volatility, working paper (2006).

[3] Mercurio, F.: Joining the SABR and Libor models together, *Risk*, 2009.

[4] Rebonato, R., McKay, K., and White, R.: *The SABR / LIBOR Market Model: Pricing, Calibration and Hedging for Complex Interest-Rate Derivatives*, Wiley (2009).