LIBOR market model and its uses

Andrew Lesniewski, Ph.D.
Managing Director
andrewl@ellington.com

Ellington Management Group
53 Forest Avenue
Old Greenwich, CT 06870

Overview

- Dynamics of Libor forward rates and martingale measures
- Volatility structures of Libor forwards
- Asymptotic solution to the LMM
- The vanilla system
- Valuation of caps and floors
- Asymptotic valuation of swaptions
- Parametrization of instantaneous volatilities and correlations
- Stable least square optimization

Overview

- Discretization of the LMM and effective methods for generating MC paths
- Accurate and fast asymptotic calculation of the drift terms
- Valuation of non-callable securities
- Valuation of securities with embedded Bermudan options
- LMM and securities with uncertain time horizons: prepayment risk, credit risk
- LMM and the SABR model of stochastic volatility
- Risk management with LMM: delta, gamma, vega, etc.

Dynamics of LMM

We consider a sequence of approximately equally spaced dates

$$0 \leq T_0 < T_1 < \ldots < T_N,$$

which are called the *standard tenors*. A standard Libor forward rate

$$L^j, \quad j = 0, 1, \ldots, N-1,$$

is associated with a forward rate agreement which starts on T_j and ends on T_{j+1}. Usually, we assume $N = 120$ and the L^j's are 3 month Libor forward rates. These dates are the actual start and end dates of the contracts rather than the fixing dates. For simplicity, we disregard this distinction. Proper implementation takes it into account.

Dynamics of LMM

We model L^j as a continuous time stochastic process $L^j(t)$, $0 \leq t \leq T_{j-1}$ (killed at $t = T_{j-1}$!). The dynamics of the forward process is driven by an N-dimensional, correlated Wiener process $W_0(t), \ldots, W_{N-1}(t)$. The probability measure associated with this Wiener process is denoted by P. We let ρ_{jk} denote the instantaneous correlation between $W_j(t)$ and $W_k(t)$, i.e.

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

To motivate the form of the stochastic differential equations describing the dynamics of the L^j's, consider the world in which there is no volatility of interest rates, i.e. for all j's $L^j(t) = L^j_0$, or

$$dL^j(t) = 0.$$

Dynamics of LMM

The fact that the rates are stochastic leads us to replace this with a system of stochastic differential equations:

$$dL^j(t) = \Delta^j(L(t), t)\, dt + C^j(L(t), t)\, dW_j(t).$$

The first term on the right hand side is called the *drift term*, and the second term is called the *diffusion term*. The no arbitrage requirement forces a relationship between the drift and the diffusion terms. The form of the drift term depends on the choice of numeraire.

Define $\gamma : [0,\ T_N] \to \mathbb{Z}$ by

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

Dynamics of LMM

We consider two types of numeraire:

- **Forward**. For $k = 1, \ldots, N$, choose as numeraire the zero coupon bond

$$P_k(t) = P_{\mathrm{stub}}(t) \prod_{\gamma(t) \le i \le k} \frac{1}{1 + \delta_i L^i(t)},$$

where $P_{\mathrm{stub}}(t)$ is the "stub adjustment" for the period $[t, T_{\gamma(t)}]$.

- **Spot**. Choose as numeraire the rolling bank account

$$P_0(t) = \frac{P_{\gamma(t)-1}(t)}{\prod_{1 \le i \le \gamma(t)-1} P_i(T_{i-1})}.$$

Let Q_k, $k = 0, 1, \ldots, N$ denote the appropriate equivalent martingale measure.

Dynamics of LMM

Under the forward measure Q$_k$, $k = 1, \ldots, N$, the dynamics of the forward rate $L^k(t)$ is a martingale. Indeed,

$$L^k(t) = \frac{1}{\delta_k} \frac{P_{k-1}(t) - P_k(t)}{P_k(t)},$$

i.e. it is the difference of two assets. Consequently, the dynamics of $L^k(t)$ should be driftless:

$$dL^k(t) = C^k\left(L^k(t), t\right) dW_k(t),$$

where $C^k\left(L^k(t), t\right)$ is an instantaneous volatility function.

Dynamics of LMM

This and the change of numeraire techniques allow us to determine the dynamics for $j \neq k$. Changing numeraire from $P_j(t)$ to $P_k(t)$, we find that:

$$dL^j(t) = C^j\left(L^j(t), t\right)$$
$$\times \begin{cases} -\sum_{j+1 \le i \le k} \dfrac{\rho_{ji}\delta_i C^i\left(L^i(t), t\right)}{1+\delta_i L^i(t)} dt + dW_j(t), & \text{if } j < k, \\ \sum_{k+1 \le i \le j} \dfrac{\rho_{ji}\delta_i C^i\left(L^i(t), t\right)}{1+\delta_i L^i(t)} dt + dW_j(t), & \text{if } j > k\,. \end{cases}$$

Dynamics of LMM

Finally, under the spot measure ($k = 0$), the LMM dynamics reads:

$$dL^j(t) = C^j\left(L^j(t), t\right)\left(\sum_{\gamma(t)\le i\le j} \frac{\rho_{ji}\delta_i C^i\left(L^i(t), t\right)}{1+\delta_i L^i(t)}\,dt + dW_j(t)\right).$$

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

$$L^j(0) = L_0^j,$$

where L_0^j is the current value of the forward which is implied by the current yield curve.

Factor reduction

In a market spanning 30 years, there are 120 quarterly Libor forwards (i.e. 120 stochastic factors). This poses severe problems with the model's implementation:

- The "curse of dimensionality" kicks in, leading to unacceptably slow performance.
- The parameters of the model are severely underdetermined and the calibration of the model becomes unstable.

We need a small number d of independent Brownian motions $Z_a(t)$, $a = 1, \ldots, d$,

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

driving the process. Typically, $d = 1, 2, 3$, or 4.

Factor reduction

We express the “true” Brownian motions in terms of the Z’s:

$$dW_j(t) = \sum_{1 \le a \le d} U_j^a \, dZ_a(t),$$

where U is an $N \times d$ matrix so that $UU' \simeq \rho$. We rewrite the dynamics of the model in terms of the independent Brownian motions:

$$dL^j(t) = \Delta^j(L(t), t)\, dt + \sum_{1 \le a \le d} B^{ja}\left(L^j(t), t\right) dZ_a(t),$$

where

$$B^{ja}\left(L^j(t), t\right) = U_j^a \, C^j\left(L^j(t), t\right).$$

Factor reduction

This system gives the *factor reduced* LMM dynamics. We interpret $\rho_d \equiv UU'$ as the correlation matrix corresponding to the factor reduced LMM dynamics. Clearly, its rank is equal to d, and

$$\rho_d = \sum_{1 \leq a \leq d}^{d} E_a E_a',$$

where E_a denotes the a-th column of the matrix U.

Low noise expansion

There is no known closed form solution to the LMM dynamics. We can construct an approximate solution by means of a *low noise expansion*: We introduce a "smallness parameter" ε in front of the diffusion coefficients. This parameter is set to 1 at the end of the computation. Since the drift coefficients are quadratic in the diffusion coefficients, we multiply them by ε^2. Expanding in powers of ε then yields:

$$
\begin{aligned}
L^j(t) = L_0^j &+ \sum_{1 \le a \le d} \int_0^t B^{ja}(L_0,\, s)\, dZ_a(s) + \int_0^t \Delta^j(L_0,\, s)\, ds \\
&+ \sum_{\substack{1 \le a,b \le d \\ 1 \le k \le N}} \iint_{0 \le u \le s \le t} B^{ka}(L_0,\, u)\, \frac{\partial B^{jb}(L_0,\, s)}{\partial L^k}\, dZ_a(u)\, dZ_b(s) \\
&+ \ldots\ .
\end{aligned}
$$

Low noise expansion

The low noise solution implies the following asymptotic expansion for the expected value of $\mathsf{E}\left[L^j(t)\right]$:

$$\mathsf{E}\left[L^j(t)\right] = L_0^j + \int_0^t \Delta^j(L_0,\, s)\, ds + \sum_{\substack{1 \leq a,b \leq d \\ 1 \leq k \leq N}} \rho_{ab} \int_0^t B^{ka}(L_0,\, s)\, \frac{\partial B^{jb}(L_0,\, s)}{\partial L^k}\, ds + \ldots\, .$$

This formula shows that the naive expected value, namely today's forward, has to be adjusted by a convexity correction which depends on the market volatility.

Low noise expansion

The covariance matrix Cov $\left[L^j(t), L^k(t)\right]$ is asymptotically given by:

$$\text{Cov}\left[L^j(t), L^k(t)\right] = \sum_{1\leq a,b\leq d} \rho_{ab} \int_0^t B^{ja}(L_0, s)\, B^{kb}(L_0, s)\, ds$$

$$+ \sum_{\substack{1\leq a,b,a',b'\leq d\\ 1\leq i,i'\leq N}} \rho_{aa'}\rho_{bb'} \int_0^t \int_0^t B^{ia}(L_0, s)\, B^{i'a'}(L_0, s')$$

$$\times \frac{\partial B^{jb}(L_0, s)}{\partial L^i} \frac{\partial B^{kb'}(L_0, s')}{\partial L^{i'}}\, ds\, ds' + \ldots\,.$$

The first term in the formula above is what one would naively expect, while the second term is a convexity correction to this quantity.

Optimization and Tikhonov regularization

In the following, we will face various multidimensional optimization problems, all of which are *ill posed* (i.e. the number of unknown parameters is smaller than the number of data points). Their solution requires minimizing a loss function $\mathcal{L}_0(x)$:

$$\mathcal{L}_0(x) = \sum_i r_i(x)^2 \,,$$

with nonlinear residuals $r_i(x)$. Such problems are typically poorly conditioned and numerically finicky. Instead, we will consider a modified loss function:

$$\mathcal{L}(x) = \mathcal{L}_0(x) + \lambda \|Sx\|^2,$$

where S is a suitable *smoothing operator*.

Vanilla system

Vanilla system consists of:

- Forward curve stripper and swap pricing model. For our purposes, the forward curve is represented as a sequence of N LIBOR forwards L_0^j, $j = 0, \ldots, N$. These forwards are used to calibrate the model, and as the initial value to evolve the LIBOR forwards.
- Cap volatility stripper and cap / floor pricing model. The output of the cap vol stripper is a sequence of N caplet volatilities. They are used to calibrate the volatility structure of the model.
- Swaption pricing model. Swaption volatilities are also used to calibrate the volatility structure of the model.

Stripping cap volatility

Caps and floors are baskets of calls and puts on Libor forwards struck at the same rate. Because of the complexity of their description (a 10 year cap involves 39 caplets!), the market quotes them in terms of a premium or a single “flat” volatility. This flat volatility has the property that, when inserted in the pricing formula, it reproduces the option premium. In reality, caplet volatility exhibits a very pronounced term structure. Typical of it is the persistence of a “volatility hump” which usually peaks somewhere between 6 months and 2 years. The process of constructing implied caplet volatility from market quotes is sometimes called *stripping*.

Stripping cap volatility

We use a two step stripping algorithm.

- We fit the caplet volatilities to the *hump function*:

$$H(t) = (\alpha + \beta t) \exp(-\lambda t) + \mu.$$

- Once α, β, λ, and μ have been calibrated, we parameterize the cap volatility curve by means of cubic B-splines and optimize to nail down the fit to market prices.

The outcome of stripping is a sequence ζ_j, $j = 0, \ldots, N-1$ of N at the money caplet volatilities. Specifically, ζ_j is the implied volatility of the caplet expiring at T_j.

Approximate valuation of swaptions

A key ingredient of efficient calibration of the LMM is rapid and accurate swaption valuation. This is hard, as:

- A swap rate is a non-linear function of the underlying Libor forward rates.
- The stochastic differential equation for the swap rate cannot be solved in closed form, and pricing swaptions requires Monte Carlo simulations.
- This poses a serious issue, as such simulations are costly.

We describe an approximation for rapid swaption valuation, based on the low noise expansion, which we use to calibrate the model.

Approximate valuation of swaptions

Since the purpose of the approximate formulas is to provide an efficient calibration tool, we focus on benchmark swaptions only. Consider a benchmark forward starting swap. The start date of the swap is denoted by T_m, and its end date is denoted by T_n. The *level function* of the swap is defined by:

$$A^{mn}(t) = \sum_{m \le j \le n-1} \alpha_j P_j(t),$$

where α_j are the day count fractions for fixed rate payments. Typically, the payment frequency on the fixed leg is not the same as that on the floating leg. This fact causes a bit of a notational nuisance but needs to be taken properly into account for accurate pricing.

Approximate valuation of swaptions

The forward swap rate is given by:

$$S^{mn}(t) = \frac{P_m(t) - P_n(t)}{A^{mn}(t)}$$
$$= \frac{1}{A^{mn}(t)} \sum_{m \le j \le n-1} \delta_j L^j(t) P_j(t) \quad .$$

Incidentally, the formula above shows that the swap rate is a martingale under the measure associated with the level function.

The swap rate process in LMM can be written as:

$$dS(t) = \Omega(L, t)\, dt + \sum_{m \le j \le n-1} \Lambda^j(L, t)\, dW_j(t)\, .$$

Approximate valuation of swaptions

Here

$$\Omega = \sum_{m \le j \le n-1} \frac{\partial S}{\partial L^j}\, \Delta^j + \frac{1}{2} \sum_{m \le j,k \le n-1} \rho_{jk}\, \frac{\partial^2 S}{\partial L^j \partial L^k}\, C^j C^k\,,$$

and

$$\Lambda^j = \frac{\partial S}{\partial L^j}\, C^j.$$

Explicitly, Λ^j is given by

$$\Lambda^j\left(L,t\right) = C^j\left(L^j,t\right)\Big(R_j\left(L,t\right) + \Xi_j\left(L,t\right)\Big),$$

where the leading term is

$$R_j\left(L,t\right) = \frac{P_j\left(t\right)}{A\left(t\right)}\,,$$

Approximate valuation of swaptions

and the correction term is given by:

$$\Xi_j = \frac{\delta_j}{1+\delta_j L^j} \left(S \sum_{j \leq l \leq n-1} \alpha_l R_l - \sum_{j \leq l \leq n-1} \delta_l L_l R_l \right).$$

To use this dynamics effectively, we 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_0^j. This amounts to “freezing” the curve at its current shape.

Approximate valuation of swaptions

Within this approximation, the coefficients are deterministic:

$$\Lambda^j (L, t) \approx \Lambda^j (L_0, t),$$
$$\Omega (L, t) \approx \Omega (L_0, t).$$

The expected value of the swap rate is given by

$$\mathsf{E}\left[S(t)\right] = S_0 + \int_0^t \Omega (L_0, s)\, ds,$$

and its variance is

$$\mathsf{Var}\left[S(t)\right] = \sum_{m \le j, j' \le n-1} \rho_{jj'} \int_0^t \Lambda^j (L_0, s)\, \Lambda^{j'} (L_0, s)\, ds.$$

Approximate valuation of swaptions

The normal volatility ζ_{mn} of a swaption with expiry T_m is

$$\zeta_{mn} = \sqrt{\frac{1}{T_m}\,\mathsf{Var}\left[S^{mn}\right]}\ .$$

Consequently, its frozen curve approximation $\zeta_{0,mn}$ is given by

$$\zeta_{0,mn}^2 = \frac{1}{T_m} \sum_{m \leq j, j' \leq n-1} \rho_{jj'} \int_0^t \Lambda^j\left(L_0, s\right) \Lambda^{j'}\left(L_0, s\right) ds\ .$$

This formula is easy to implement in code, and leads to remarkably accurate results. One could, of course, go beyond this approximation, at the expense of producing a complicating and increasingly unwieldy analytic expressions.

Structure of instantaneous correlations

An important input into the model calibration is determining the instantaneous correlation matrix $\rho = \{\rho_{jk}\}_{0 \leq j,k \leq N-1}$. The dimensionality of ρ is $N(N+1)/2$, posing an issue of finding a stable calibration procedure. Possible strategies include:

- **Semi-definite programming**. Try to imply the correlations from the cap / floor and swaption markets. This approach leads to non-intuitive results and is prone to overfitting.
- **Historical data**. This approach, in conjunction with the *principal component analysis* leads to stable correlation structure. It may not reflect the current market correlations.
- **Parameterized correlations**. May be the best choice.

Structure of instantaneous correlations

A realistic and flexible parametric form of instantaneous correlations is given by

$$\rho_{jk} = \rho_\infty + (1 - \rho_\infty) \exp\left(-\lambda \, \frac{|T_j - T_k|}{1 + \kappa \min(T_j, T_k)} \right).$$

Caution: this parametrization produces a matrix that is only approximately positive definite. It has some clear advantages:

- It is intuitive: ρ_∞ measures the overall level of correlations, λ is the decay rate of correlations, and κ describes the short end decorrelation.
- It is easy to calibrate: only three parameters are involved.
- Perturbing the parameters $\rho_\infty, \lambda, \kappa$ leads to meaningful risk measures.

Parametrization of instantaneous volatilities

So far we have been working with a general instantaneous volatility $C^j\left(L^j(t), t\right)$ for the forward $L^j(t)$. In the implementation, we assume $C^j\left(L^j(t), t\right)$ to be one of the following standard models:

$$C^j\left(L^j(t), t\right) = \begin{cases} \sigma^j(t) & \text{(normal model)}, \\ \sigma^j(t)\, L^j(t)^{\beta_j} & \text{(CEV model)}, \\ \sigma^j(t)\, L^j(t) & \text{(lognormal model)}, \\ \sigma^j(t)\, L^j(t) + \delta^j & \text{(shifted lognormal model)}, \end{cases}$$

where the functions $\sigma^j(t)$ are deterministic, and where $0 \leq \beta^j \leq 1$, $\delta^j \geq 0$. These functions have to be suitably regularized at $L^j \to 0$ to prevent negative rates.

Parametrization of instantaneous volatilities

For the purpose of calibration we require that the deterministic volatility components $\sigma^j(t)$ are piecewise constant. That leads to the following parametrization of the instantaneous volatility:

$\sigma^j(t) \diagdown t \in$	$[T_0, T_1)$	$[T_1, T_2)$	$[T_2, T_3)$	$\ldots$	$[T_{N-1}, T_N)$
$\sigma^0(t)$	0	0	0	$\ldots$	0
$\sigma^1(t)$	$\sigma_{1,0}$	0	0	$\ldots$	0
$\sigma^2(t)$	$\sigma_{2,0}$	$\sigma_{2,1}$	0	$\ldots$	0
$\sigma^3(t)$	$\sigma_{3,0}$	$\sigma_{3,1}$	$\sigma_{3,2}$	$\ldots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\sigma^{N-1}(t)$	$\sigma_{N-1,0}$	$\sigma_{N-1,1}$	$\sigma_{N-1,2}$	$\ldots$	0

Parametrization of instantaneous volatilities

The table above contains 7140 parameters (assuming $N = 120$), and the problem is vastly *overparametrized*!

A natural remedy to the overparameterization problem is to assume that the instantaneous volatility is *stationary*, i.e.,

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

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.

Parametrization of instantaneous volatilities

With the stationary volatility assumption we have the following parametrization of the instantaneous volatility structure:

$\sigma^j(t) \diagdown t \in$	$[T_0, T_1)$	$[T_1, T_2)$	$[T_2, T_3)$	$\ldots$	$[T_{N-1}, T_N)$
$\sigma^0(t)$	0	0	0	$\ldots$	0
$\sigma^1(t)$	σ_1	0	0	$\ldots$	0
$\sigma^2(t)$	σ_2	σ_1	0	$\ldots$	0
$\sigma^3(t)$	σ_3	σ_2	σ_1	$\ldots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\sigma^{N-1}(t)$	σ_{N-1}	σ_{N-2}	σ_{N-3}	$\ldots$	0

Parametrization of instantaneous volatilities

This assumption is unsuitable for accurate calibration of the model. The financial reason is the phenomenon of *mean reversion* of long term rates. One cannot take it into account by adding an Ornstein-Uhlenbeck style drift term to the LMM dynamics as this would violate the arbitrage freeness of the model. One can, however, achieve a similar effect by modifying the instantaneous volatility function. We introduce a *volatility kernel* function $K(\tau, \lambda)$ to account for the deviation from the purely stationary model. A convenient form of the volatility kernel is

$$K(\tau, \lambda) = \exp(-\lambda\tau).$$

For each maturity T_j we choose a parameter λ_j, and set

$$K_{j,i} = K(T_j - T_i, \lambda_j).$$

Parametrization of instantaneous volatilities

We are thus led to assume the following structure of the instantaneous volatility:

$\sigma^j(t) \diagdown t \in$	$[T_0, T_1)$	$[T_1, T_2)$	$\ldots$	$[T_{N-1}, T_N)$
$\sigma^0(t)$	0	0	$\ldots$	0
$\sigma^1(t)$	$\sigma_1 K_{1,0}$	0	$\ldots$	0
$\sigma^2(t)$	$\sigma_2 K_{2,0}$	$\sigma_1 K_{2,1}$	$\ldots$	0
$\sigma^3(t)$	$\sigma_3 K_{3,0}$	$\sigma_2 K_{3,1}$	$\ldots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\sigma^{N-1}(t)$	$\sigma_{N-1} K_{N-1,0}$	$\sigma_{N-2} K_{N-1,1}$	$\ldots$	0

Parametrization of instantaneous volatilities

The lower triangular matrix above, LMM's internal representation of volatility, is referred to as the *LMM volatility surface*.

The structure above may still overparameterize the instantaneous volatility, depending on the number of benchmark options used to calibrate the model. To further reduce the number of parameters we express the σ_j's and λ_j's as linear interpolations of a smaller number of auxiliary parameters.

Least square 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. The at the money volatility of the caplet expiring at T_m is given by:

$$\begin{aligned}\zeta_m\left(\sigma_1,\ldots,\sigma_m,\lambda_m\right)^2 &= \frac{1}{T_m}\sum_{0\leq i\leq m-1}\sigma_{m-i}^2\int_{T_i}^{T_{i+1}} K\left(T_m-t,\lambda_m\right)^2 dt\\ &\approx \frac{1}{T_m}\sum_{0\leq i\leq m-1}\sigma_{m-i}^2 K_{m,i}^2\left(T_{i+1}-T_i\right).\end{aligned}$$

Least square optimization

The at the money volatility of the swaption expiring at T_m into a swap maturing at T_n is approximately equal to

$$\zeta_{m,n}\left(\sigma, \lambda_m, \ldots, \lambda_{n-1}\right)^2 = \frac{1}{T_m} \sum_{m \leq j,k \leq n-1} \rho_{jk} \Lambda^j_{m,n} \Lambda^k_{m,n} \times \sum_{i=0}^{m-1} \sigma_{j-i} \sigma_{k-i} K_{j,i} K_{k,i} \left(T_{i+1} - T_i\right).$$

Here, Λ^j are simply rescaled versions of the corresponding functions which were defined and calculated asymptotically.

Least square optimization

We calibrate the model by means of the Levenberg-Marquardt optimization. The loss function is given by:

$$\begin{aligned}\mathcal{L}(\sigma,\lambda) = {} & \frac{1}{2}\sum_m w_m \left(\zeta_m(\sigma,\lambda) - \overline{\zeta}_m\right)^2 \\ & + \frac{1}{2}\sum_{m,n} w_{m,n}\left(\zeta_{m,n}(\sigma,\lambda) - \overline{\zeta}_{m,n}\right)^2 \\ & + \frac{1}{2}\alpha\sum_j (\Delta\sigma)_j^2\,,\end{aligned}$$

where $\overline{\zeta}_m$ and $\overline{\zeta}_{m,n}$ are the market observed caplet and swaption volatilities. The coefficients w_m and $w_{m,n}$ are weights which allow one select the accuracy of calibration of each of the instruments. The last term is a Tikhonov regularizer.

Approximation schemes for LMM

For the purpose of asset pricing we solve numerically the SDEs defining the LMM dynamics. Such a solution is a collection of Monte Carlo paths, each of which represents a future forward rate scenario.

We choose a sequence of *event dates* $t_0, t_1, \ldots, t_m$, and denote by $L_n^j \simeq L^j(t_n)$ the approximate solution. We also use

$$\Delta_n^j = \Delta^j (L_n, t_n) ,$$

$$C_n^{ja} = C^{ja} (L_n, t_n) ,$$

to denote the values of the drift and diffusion terms at time t_n.

Approximation schemes for LMM

The simplest numerical scheme is **Euler's scheme**. It consists in replacing the differentials by finite differences. For $n = 1, \ldots, m$, we let

$$\delta t_n = t_n - t_{n-1},$$
$$dW_{na} = W_a(t_n) - W_a(t_{n-1}),$$

denote the time and Wiener process increments,respectively. Note that $dW_{na} \sim N\left(0, \sqrt{\delta t_n}\right)$. Then, Euler's solution scheme reads:

$$L_{n+1}^j = L_n^j + \Delta_n^j \delta t_n + \sum_{1 \le a \le d} U_j^a C_n^j \, dW_{na} \,.$$

Euler's scheme converges at the rate of $\max \sqrt{\delta t_n}$, as $m \to \infty$.

Approximation schemes for LMM

A somewhat more accurate method is **Milstein's scheme**. It is a refinement of Euler's scheme which converges at the rate of $\max \delta t_n$, as $m \to \infty$. We set

$$\Upsilon_n^{jab} = U_j^a U_j^b C^j \left(L_n^j,\, t_n\right) \frac{\partial C^j}{\partial L^j} \left(L_n^j,\, t_n\right) .$$

Then, Milstein's scheme for the LMM reads:

$$L_{n+1}^j = L_n^j + \left(\Delta_n^j - \frac{1}{2} \sum_{1 \le a \le d} \Upsilon_n^{jaa} \right) \delta t_n$$
$$+ \sum_{1 \le a \le d} U_j^a C_n^j \, dW_{na} + \frac{1}{2} \sum_{1 \le a,b \le d} \Upsilon_n^{jab} \, dW_{na} \, dW_{nb} \, .$$

Approximation schemes for LMM

Having discretized the dynamics, we have to select a method for simulating the Wiener process. An efficient method relies on the *spectral decomposition* of the covariance matrix of $W(t)$ sampled at $t_0, t_1, \ldots, t_m$. The covariance matrix is explicitly given by:

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

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

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

with orthonormal eigenvectors E_j's and eigenvalues

$$\lambda_0 \geq \ldots \geq \lambda_m \geq 0.$$

Approximation schemes for LMM

We let $E_j(t_n)$ denote the n-th component of E_j, and represent the discretized Wiener process as

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

where ξ_j are 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. In practice, we may want to use only a certain portion of the spectral representation 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 is lower accuracy.

Computing the drift terms

Evaluating the drift terms along each Monte Carlo path is time consuming and accounts for over 50% of total simulation time. Typically they are small compare to the initial values of the Libor forwards, and it is desirable to develop an efficient methodology for accurate approximate evaluation of the drift terms.
The first and simplest approach consist in “freezing” the values of $L^j(t)$ at the initial value $L_0^j \equiv L^j(0)$. We precompute the values

$$\Delta_0^j \equiv \Delta^j(L_0, 0),$$

and use them for the drift terms throughout the simulation. This approximation, the *frozen curve approximation*, is rather crude, and does not perform very well for long dated instruments.

Computing the drift terms

The *order* $3/4$ *approximation*, uses the next order term in the low noise expansion:

$$\Delta^j_{1/2}(t) \equiv \Delta^j_0(L_0, 0) + \Gamma^{a,j}(L_0, 0)\, Z_a(t) + \Omega^j(L_0, 0)\, t.$$

Under the forward measure Q_k, the coefficients $\Gamma^{a,j}$ are given by:

$$\Gamma^{a,j} = \begin{cases} -C^j \displaystyle\sum_{j+1 \le i \le k} \frac{\rho_{ji}\delta_i C^i}{1+\delta_i L^i} \left[U^a_j \frac{\partial C^j}{\partial L^j} + U^a_i \left(\frac{\partial C^i}{\partial L^i} - \frac{\delta_i C^i}{1+\delta_i L^i} \right) \right], & \text{if } j < k, \\ 0, & \text{if } j = k, \\ C^j \displaystyle\sum_{k+1 \le i \le j} \frac{\rho_{ji}\delta_i C^i}{1+\delta_i L^i} \left[U^a_j \frac{\partial C^j}{\partial L^j} + U^a_i \left(\frac{\partial C^i}{\partial L^i} - \frac{\delta_i C^i}{1+\delta_i L^i} \right) \right], & \text{if } j > k, \end{cases}$$

Computing the drift terms

and the coefficients Ω^j are given by:

$$\Omega^j = \begin{cases} -\sum_{j+1\le i\le k} \frac{\rho_{ji}\delta_i}{1+\delta_i L^i}\left[\Delta_0^j C^i \frac{\partial C^j}{\partial L^j} + \Delta_0^i C^j \left(\frac{\partial C^i}{\partial L^i} - \frac{\delta_i C^i}{1+\delta_i L^i}\right)\right], & \text{if } j < k, \\ 0, & \text{if } j = k, \\ \sum_{k+1\le i\le j} \frac{\rho_{ji}\delta_i}{1+\delta_i L^i}\left[\Delta_0^j C^i \frac{\partial C^j}{\partial L^j} + \Delta_0^i C^j \left(\frac{\partial C^i}{\partial L^i} - \frac{\delta_i C^i}{1+\delta_i L^i}\right)\right], & \text{if } j > k. \end{cases}$$

Under the spot measure,

$$\Gamma^{a,j} = C^j \sum_{\gamma(t)\le i\le j} \frac{\rho_{ji}\delta_i C^i}{1+\delta_i L^i}\left[U_j^a \frac{\partial C^j}{\partial L^j} + U_i^a \left(\frac{\partial C^i}{\partial L^i} - \frac{\delta_i C^i}{1+\delta_i L^i}\right)\right],$$

Computing the drift terms

and

$$\Omega^j = \sum_{\gamma(t) \leq i \leq j} \frac{\rho_{ji}\delta_i}{1+\delta_i L^i} \left[\Delta_0^j C^i \frac{\partial C^j}{\partial L^j} + \Delta_0^i C^j \left(\frac{\partial C^i}{\partial L^i} - \frac{\delta_i C^i}{1+\delta_i L^i} \right) \right].$$

The order $3/4$ approximation leads to excellent accuracy.

On might easily refine this approach by computing terms of higher order in stochastic Taylor's expansion. This leads, however, to more complex and computationally expensive formulas, and the benefit of using an asymptotic expansion disappears. The order $1/2$ approximation appears to offer the best performance versus accuracy profile.

Valuation of non-callable securities

Valuations within the LMM are based on the arbitrage free pricing law. For our purposes, the equivalent martingale measure Q is either one of the forward measures or the spot measure. The conditional expected value is calculated by means of Monte Carlo simulations. A convenient choice of the numeraire for all valuations is the spot numeraire. The sample paths generation method described above produces low variance estimates of the expected value. In addition, one might use a generic variance reduction method (such as antithetic variables) in order to further reduce the sampling variance. For most instruments, a relatively small number of Monte Carlo paths leads to accurate and stable valuations. As few as 1000 paths are sufficient to produce reliable prices.

Valuation of Bermudan options

Valuation of securities with embedded American (or Bermudan) options is more difficult, as the traditional Monte Carlo approach is inefficient. The problem is that while optimal exercise of a Bermudan option requires solving a backward induction problem, Monte Carlo paths evolve forward in time. Recently, a number of efficient approximate Monte Carlo algorithms for pricing Bermudan and American options have been proposed. The approach we outline is inspired by the Longstaff - Schwarz approach. It is based on a sequence of nested *chaos expansions* of the continuation values at each exercise date. As few as 5000 Monte Carlo paths lead to reliable pricing.

Valuation of Bermudan options

We fix a measure Q_k and consider a Bermudan swaption which can be exercised on dates $t_1, \ldots, t_M$. If exercised on t_j, it allows the holder to enter into a swap starting on (or about) t_j. We let $p(t) = p(L(t), t)$ denote the payoff function. For a stopping time τ, we let $p^\tau(t) = p(t \wedge \tau)$. Then its value at time t_0 (today) equals

$$V(t_0) = \sup_{\tau \in \mathcal{T}(t_1, \ldots, t_M)} \mathsf{E}\left[p^\tau \mid \mathcal{F}_{t_0}\right],$$

where $\mathcal{T}(t_1, \ldots, t_M)$ denotes the set of stopping times taking values in the set $\{t_1, \ldots, t_M\}$. The stopping time at which the maximum is attained is called the *optimal stopping time*.

Valuation of Bermudan options

In order to simplify the notation, we assume one factor only, $d = 1$ (nothing gets lost). We introduce the notation:

$$X_j = \frac{W(t_j)}{\sqrt{t_j}}, \quad j = 1, \ldots, M,$$

so that $X_j \sim N(0,1)$. From the properties of a Wiener process,

$$X_j = \sqrt{\frac{t_{j-1}}{t_j}}\, X_{j-1} + \sqrt{1 - \frac{t_{j-1}}{t_j}}\, \xi_j, \;\; \text{with} \;\; \xi_j \sim N(0,1).$$

We denote by $g_j(X_j) = p(L(t_j), t_j)$ the payoff function at time t_j, i.e. the present value (at time t_j) of the corresponding swap.

Valuation of Bermudan options

Let $V_j(X_j)$ denote the value of the option at time t_j. Then,

$$V_M(X_M) = g_M(X_M).$$

For $j = M-1, \ldots, 1$, the value of the option is the greater of the immediate payoff and the *continuation value* of the option:

$$V_j(X_j) = \max\left(g_j(X_j),\; D(t_j, t_{j+1}) \mathsf{E}\left[V_{j+1}(X_{j+1}) | X_j\right]\right),$$

where $D(t_j, t_{j+1}) = P_k(t_{j+1}) / P_k(t_j)$. Today's value of the option is

$$V_0(X_0) = D(t_0, t_1) \mathsf{E}\left[V_1(X_1) | X_0\right].$$

Valuation of Bermudan options

The continuation value of the option at time t_j is

$$C_j\left(X_j\right) = D\left(t_j, t_{j+1}\right) \mathsf{E}\left[V_{j+1}\left(X_{j+1}\right) | X_j\right],$$

i.e.

$$V_j\left(X_j\right) = \max\left(g_j\left(X_j\right),\ C_j\left(X_j\right)\right).$$

We also have a recursion for the continuation values:

$$\begin{aligned} C_M\left(X_M\right) &= 0, \\ C_j\left(X_j\right) &= D\left(t_j, t_{j+1}\right) \mathsf{E}\left[\max\left(g_{j+1}\left(X_{j+1}\right),\ C_{j+1}\left(X_{j+1}\right)\right) | X_j\right], \end{aligned}$$

for $j = M-1, \ldots, 1$. Today's option price is

$$V_0\left(X_0\right) = C_0\left(X_0\right).$$

Valuation of Bermudan options

We now would like to write

$$V_j\left(X_j\right) = \sum_{0 \le k < \infty} a_{jk}\, H_k\left(X_j\right),$$

where $H_k\left(X_k\right)$ denotes the normalized Hermite polynomial of degree k. The Fourier coefficients a_n are calculated by

$$a_{jk} = \int_{-\infty}^{\infty} V_j\left(X\right) H_k\left(X\right) d\mu\left(X\right).$$

The martingale property of Hermite polynomials yields:

$$\mathsf{E}\left[V_j\left(X_j\right) | X_{j-1}\right] = \sum_{0 \le k < \infty} \left(\frac{t_{j-1} - t_0}{t_j - t_0}\right)^{k/2} a_{jk}\, H_k\left(X_{j-1}\right).$$

Valuation of Bermudan options

The expansions with respect to the Hermite polynomials are truncated at some finite order κ, i.e.:

$$V_j\left(X_j\right) \simeq \sum_{0 \le k \le \kappa} a_{jk}\, H_k\left(X_j\right),$$

$$\mathsf{E}\left[V_j\left(X_j\right) \middle| X_{j-1}\right] \simeq \sum_{0 \le k \le \kappa} \left(\frac{t_{j-1} - t_0}{t_j - t_0}\right)^{k/2} a_{jk}\, H_k\left(X_{j-1}\right).$$

Relatively low values of κ, $\kappa \sim 5$, lead to good numerical results.

Valuation of Bermudan options

The Fourier coefficients a_k are replaced by their Monte Carlo estimates. We shall choose the coefficients so as to minimize the square error:

$$\frac{1}{2} \sum_{1 \le i \le N} \left(V_j \left(X_j^{(i)} \right) - \sum_{0 \le k \le \kappa} a_{jk} H_k \left(X_j^{(i)} \right) \right)^2 .$$

This leads to the following estimate for the Fourier coefficients:

$$a_j \simeq G_j^{-1} v_j ,$$

where G_j is a matrix whose components are:

Valuation of Bermudan options

$$(G_j)_{kl} = \frac{1}{N} \sum_{1 \leq i \leq N} H_k\left(X_j^{(i)}\right) H_l\left(X_j^{(i)}\right),$$

and where v_j is a vector with the components:

$$(v_j)_k = \frac{1}{N} \sum_{1 \leq i \leq N} V_j\left(X_j^{(i)}\right) H_k\left(X_j^{(i)}\right).$$

From the performance point of view, is worthwhile to accelerate the convergence of this series. Payoff functions of options are, typically, not smooth. This causes slowdowns of the rate of convergence at the points where the payoff has kinks (as in the usual calls or puts) or discontinuities (as in digital options).

Valuation of Bermudan options

Having estimated the conditional expected values (and thus calculated the consecutive continuation values), we construct the optimal stopping time as follows.

- Along a Monte Carlo path, find the earliest of the dates t_j, where the immediate exercise outweighs the continuation value.
- Assign this date to the optimal stopping time along the path.
- Calculate the average over all Monte Carlo paths.

Securities with uncertain horizons

It is important to extent LMM so that instruments with risks other than interest rate risk can be valued. Certain classes of such instruments, such as mortgage backed securities (MBS) or credit sensitive instruments are characterized by the uncertainty of their time horizon. It is caused by an *event* such as *prepayment* in case of MBS, or *default* in case of credit sensitive instruments. Unlike the callable LIBOR exotics, where the termination decision is made optimally and consistently with the rates process, these risks are largely exogenous to the rates process. Below we formulate some basic assumptions about modeling these instruments.

Securities with uncertain horizons

In order to make the framework consistent with LMM we make the following assumptions.

- **Assumption 1**. The underlying uncertainty is modeled by the LIBOR market model.
- **Assumption 2**. Additional uncertainty (prepayment, default, ...) is modeled by a random time T. T is not required to be a stopping time with respect to $\mathcal{F}_t$. However, we assume that

$$\mathsf{P}\left(\{T > t\} \,|\mathcal{F}_\infty\right) = \mathsf{P}\left(\{T > t\} \,|\mathcal{F}_t\right).$$

The second assumption means that the probability of event occurring depends only on the information up to time t

Securities with uncertain horizons

- **Assumption 3**. There exist a process $\lambda(t) \geq 0$, the *intensity process*, so that

$$\mathsf{P}\left(\{T > t\} \,|\mathcal{F}_t\right) = \exp\left(-\int_0^t \lambda(s)\, ds\right).$$

  In the context of prepayment, $\lambda(t)$ is the (continuous time version of) SMM. In the context of credit, it is the hazard rate.

- **Assumption 4**. The event will occur almost surely,

$$\mathsf{P}\left(\{T < \infty\}\right) = 1.$$

Securities with uncertain horizons

Valuation of securities with uncertain time horizons is done by means of Monte Carlo simulations. The stochastic process

$$S(t) = \exp\left(-\int_0^t \lambda(s)\, ds\right)$$

is called the *survival probability*, while

$$F(t) = 1 - S(t)$$

is the *event probability*.

Securities with uncertain horizons

A cashflow of a security with event risk is given by:

$$\sum_j \Big(c_j P\left(t_j\right) S\left(t_j\right) + r_j P\left(t_j\right) \left(F_j\left(t\right) - F\left(t_{j-1}\right)\right) \Big),$$

where c_j are the known cash amounts, and r_j are the recovery rates in case an event occurs. The intensity process defining the survival probability is modeled outside of LMM. The price of the security is calculated by taking the appropriate expected value of the cashflows above.

LMM and the SABR model

The **SABR model** describes the dynamics of a single forward with stochastic volatility. It's dynamics is given by:

$$dL(t) = \sigma C(L(t))\, dW(t),$$
$$d\sigma(t) = \alpha\sigma(t)\, dZ(t),$$

where

$$\mathsf{E}[dW(t)\, dZ(t)] = r\, dt.$$

One usually chooses $C(L)$ to be of the CEV type

$$C(L) = L^{\beta}.$$

SABR captures volatility smile on vanilla options. In order to model the smile of more complex instruments, we need to suitably extend LMM.

LMM and the SABR model

We consider an extension of the LIBOR market model with stochastic volatility parameters denoted by $\sigma_t^1, \ldots, \sigma_t^N$:

$$dL^j(t) = \Delta^j(L(t), \sigma(t), t)dt + C^j(L^j(t), \sigma^j(t), t)dW^j(t),$$

$$d\sigma^j(t) = \Gamma^j(L(t), \sigma(t), t)dt + D^j(L^j(t), \sigma^j(t), t)dZ^j(t),$$

with

$$\mathsf{E}[dW^i(t)\ dW^j(t)] = \rho_{ij}dt,$$

$$\mathsf{E}[dW^i(t)\ dZ^j(t)] = r_{ij}dt,$$

$$\mathsf{E}[dZ^i(t)\ dZ^j(t)] = \pi_{ij}dt.$$

Choose the drifts so that this dynamics is arbitrage free!

LMM and the SABR model

Under the forward measure Q_k, the arbitrage free dynamics is:

$$dL_t^j = C^j(L_t^j, \sigma_t^j, t) \times \begin{cases} -\sum_{j+1\le i\le k} \dfrac{\rho_{ji}\delta_i C^i(L_t^i, \sigma_t^i, t)}{1+\delta_i L_t^i}\, dt + dW_t^j, & \text{if } j < k, \\ dW_t^j, & \text{if } j = k, \\ \sum_{k+1\le i\le j} \dfrac{\rho_{ji}\delta_i C^i(L_t^i, \sigma_t^i, t)}{1+\delta_i L_t^i}\, dt + dW_t^j, & \text{if } j > k, \end{cases}$$

$$d\sigma_t^j = D^j(L_t^j, \sigma_t^j, t) \times \begin{cases} -\sum_{j+1\le i\le k} \dfrac{r_{ji}\delta_i C^i(L_t^i, \sigma_t^i, t)}{1+\delta_i L_t^i}\, dt + dZ_t^j, & \text{if } j < k, \\ dZ_t^j, & \text{if } j = k, \\ \sum_{k+1\le i\le j} \dfrac{r_{ji}\delta_i C^i(L_t^i, \sigma_t^i, t)}{1+\delta_i L_t^i}\, dt + dZ_t^j, & \text{if } j > k\,. \end{cases}$$

LMM and the SABR model

Under the spot measure, the dynamics reads:

$$
\begin{aligned}
dL^j(t) &= C^j\left(L^j(t), \sigma^j(t), t\right) \sum_{\gamma(t) \le i \le j} \frac{\rho_{ji}\delta_i C^i\left(L^i(t), \sigma^i(t), t\right)}{1 + \delta_i L^i(t)}\, dt \\
&\quad + C^j\left(L^j(t), \sigma^j(t), t\right) dW_j(t)\ , \\
d\sigma^j(t) &= D^j\left(\sigma^j(t), t\right) \sum_{\gamma(t) \le i \le j} \frac{\rho_{ji}\delta_i C^i\left(L^i(t), \sigma^i(t), t\right)}{1 + \delta_i L^i(t)}\, dt \\
&\quad + D^j\left(\sigma^j(t), t\right) dZ_j(t)\,.
\end{aligned}
$$

Note that the model involves a large number of parameters. Judicious choices have to be made in order to produce a stable framework.

Risk management: delta

This is the most important risk factor (and the easiest to hedge).

- Choose a “hedging portfolio” consisting of vanilla instruments such as swaps, Eurodollar futures, forward rate agreements, etc:

$$\Pi_{\mathrm{hedge}} = \{B_1, \ldots, B_n\} \ .$$

- Let $\mathcal{C}_0$ denote the current forward curve (the “base scenario”). Choose a number of new micro scenarios

$$\mathcal{C}_1, \ldots, \mathcal{C}_p$$

  by perturbing a segment of $\mathcal{C}_0$. For example, $\mathcal{C}_1$ could result from $\mathcal{C}_0$ by shifting the first 3 month segment down by 1 bp.

Risk management: delta

- The vector $\delta\Pi$ of portfolio's *sensitivities* to these scenarios is

$$\delta_i \Pi = \Pi\left(\mathcal{C}_i\right) - \Pi\left(\mathcal{C}_0\right), \qquad i = 1, \ldots, p,$$

  where by $\Pi\left(\mathcal{C}_i\right)$ we denote the value of the portfolio given the forward curve $\mathcal{C}_i$.

- The matrix δB of sensitivities of the hedging instruments to these scenarios is

$$\delta_i B_j = B_j\left(\mathcal{C}_i\right) - B_j\left(\mathcal{C}_0\right).$$

  To avoid accidental colinearities between its rows or columns, one should always use more scenario than hedging instruments.

Risk management: delta

- The vector Δ of *hedge ratios* is calculated by minimizing

$$\mathcal{L}\left(\Delta\right) = \frac{1}{2}\left\|\delta B\,\Delta - \delta\Pi\right\|^2 + \frac{1}{2}\lambda\|Q\,\Delta\|^2.$$

Here, λ is an appropriately chosen small smoothness parameter, and Q is the smoothing operator. Explicitly,

$$\Delta = \left(\left(\delta B\right)^t\,\delta B + \lambda Q^t\,Q\right)^{-1}\left(\delta B\right)^t\,\delta\Pi.$$

One can think of the component Δ_j as the sensitivity of the portfolio to the hedging instrument B_j. This method of calculating portfolio sensitivities is called the *ridge regression*.

Risk management: gamma

The *gamma* of a portfolio is sometimes calculated as its global convexity characteristic. This is a rather crude measure, as portfolios typically exhibit complex convexity behaviors. A better way is to construct the portfolio gamma as the change in its delta under specified macro scenarios:

$$\Xi_0, \Xi_1, \ldots, \Xi_r,$$

with Ξ_0 base scenario (no change in rates). For example:

Ξ_{+50}	All rates up 50 basis points.
Ξ_{+25}	All rates up 25 basis points.
Ξ_{-25}	All rates down 25 basis points.
Ξ_{-50}	All rates down 50 basis points.

Risk management: gamma

For each of the macro scenarios, we calculate the deltas

$$\Delta_1, \ldots, \Delta_r.$$

The quantities:

$$\Gamma_1 = \Delta_1 - \Delta_0,$$
$$\vdots$$
$$\Gamma_r = \Delta_r - \Delta_0,$$

are the portfolio gammas under the corresponding scenarios. For intermediate market moves, the portfolio gamma can be calculated by linearly interpolating the macro scenarios.

Risk management: vega

In order to quantify the vega risk we have to first design appropriate volatility scenarios.

- **Perturb vol inputs**. Shift selected swaption and/or cap volatilities.
- **Perturb internal parameters**. As explained before, LMM builds its internal "volatility surface" $\mathfrak{S}$. We construct volatility micro scenarios by accessing $\mathfrak{S}$ and shifting selected non-overlapping segments. Let us call these scenarios

$$\mathfrak{S}_0, \mathfrak{S}_1, \ldots, \mathfrak{S}_q,$$

with $\mathfrak{S}_0 = \mathfrak{S}$ being the base scenario.

Risk management: vega

- We choose a hedging portfolio Π_{hedge} which may consist of liquid instruments such as swaptions, caps and floors, Eurodollar options, or other instruments (Bermudan options?).
- The rest is a *verbatim* repeat of the delta story. We calculate the sensitivities of the portfolio to the volatility scenarios. We calculate the sensitivities of the hedging portfolio to the volatility scenarios. Finally, we use ridge regression to find the hedge ratios.

This method of managing the vega risk works remarkably well and allows one, in particular, to separate the exposure to swaptions from the exposure to caps / floors.