Cosine densities approximations: Applications to swaptions pricing

Approximating densities to speed up pricing

Sophian Mehalla
Lucien Morice

Economic Scenario Generators (ESGs) are essential tools for insurance companies. The production of market-consistent scenarios requires the models to be calibrated with the current market information. Within ESGs, interest rate models focus the attention of practitioners. Their complexity has significantly increased over the last decade, and so has the need for fast and accurate pricing methods for derivatives. This paper describes an efficient swaptions pricing method based on density approximation with Fourier series under the LIBOR[1] Market Model with Displaced Diffusion and Stochastic Volatility (DD-SV-LMM) framework. A comparison to standard methods is made.

Motivations

ESGs are the cornerstone of many processes in insurance companies such as the computation of risk management indicators, or the valuation of long-term commitments with optional warranties depending on the economic situation. An ESG is defined as a set of models used to project the joint behaviour of relevant economic or financial risk factors over multiple scenarios. In order to be consistent with current economic conditions, simulations should be generated by models calibrated to current market data: this is the so-called *market-consistency*. The Milliman Economic Scenario Generator[2], that has been employed to lead the present work, is an ESG used worldwide.

In the case of insurance companies, modelling interest rate risk is a priority because life insurance policies consist in long-term commitments embedding optional warranties that may be activated depending on policyholder behaviour in various economic environments. Moreover, bonds and interest rate derivatives constitute a major part of the asset allocation of insurance companies. This is why, for the sake of consistency, interest rate models are usually calibrated to European swaptions.

Most interest-rate models fall into two categories: "short rate" model or "market" model. The former focusses on modelling the short rate, an unobservable factor corresponding to the return of an investment over an infinitesimal period, as in the Hull & White and G2++ models.

The LIBOR Market Model (LMM) and its different versions, on the other hand, model quantities that are directly observable in the market. In their simplest specification, they provide a theoretical framework consistent with the use of the Black or Bachelier formulas. Adding a shift coefficient to "displace" the simulated distribution into the non-positive region (i.e., allowing it to generate negative rate environments) and stochastic volatility (to better match the observed skew in the market) yields the Displaced Diffusion with Stochastic Volatility LMM (DD-SV-LMM).

In its original specification, the DD-SV-LMM contains too much randomness—roughly speaking—to be analytically tractable. To get an exploitable version of the model, some stochastic quantities are therefore "frozen" to their initial values to remove some hazard and thus simplify the model. This is the so-called freezing technique. This assumption yields a Heston-like model under which the moment-generating function can be numerically computed, and the European options can then be evaluated through moment-generating function integration. This method requires a significant computation time because the moment-generating function must be computed multiple times for each option to be priced.

In the DD-SV-LMM, one has access to the (approximate) moment-generating function of the swap rate process, which allows us, by applying Fourier transformation, to recover the density function of the process. However, this transformation is costly from a numerical point of view as it requires numerical approximation of integrals based on Gaussian quadrature. Other methods have been proposed to efficiently approximate

[1] The name of the model has been settled several years ago, motivated by use of the London Inter-Bank Offered Rate (LIBOR) as reference. Names of the models may be adapted in the future because LIBOR will no longer be published after June 2023.

[2] See https://www.milliman.com/products/economic-scenario-generator.

the density function based on polynomial expansions (see notably [MEH21]). Those methods are competitive (both in terms of computational time and data replication accuracy) but their validity domain can be restricted in some cases. The so-called "cosine expansion" method introduced in [FAN10] and presented in this document still takes advantage of the link between moment-generating and density functions but allows us, by applying a number of approximations, to perform computations without requiring cumbersome quadratures. In the end, we would be able to compute swaption prices based on this competitive approach. This method is aimed at reducing the computational time dedicated to each option price without excessive accuracy losses.

In the end, this method contributes towards proposing efficient methods for the calibration of models forming the Milliman Economic Scenarios Generator.

The DD-SV-LMM

The basic interest rate product to introduce is the zero-coupon (ZC) bond: it is a contract that delivers one unit of currency at a future date T whose price at time t is $P(t,T)$.

Let $T_j < T_{j+1}$ be two dates. A forward rate agreement (FRA) is a contract in which two counterparties agree to exchange a fixed rate against a floating one for the accrual period between T_j and T_{j+1}. At any time t prior to the beginning of the period ($t \leq T_j$), there is a unique fixed rate that makes this contract arbitrage-free. This rate corresponds to the forward rate $F_j(t)$ and it can be expressed in terms of ZC bonds as:

$$F_j(t) = \frac{1}{T_{j+1}-T_j}\left(\frac{P(t,T_j)}{P(t,T_{j+1})} - 1\right).$$

An interest rates swap (IRS) is a sequence of (off-market) FRAs over multiple time periods. Let $T_m < \cdots < T_n$ be a set of dates. An IRS corresponds to the exchange of a fixed and constant rate $S(t)$ for a fluctuating rate on the corresponding period. We have followed the convention in which the payments are made on dates $\{T_{m+1}, \ldots, T_n\}$. With the same arguments as for the FRA, the forward swap rate value at time $t \leq T_m$ can be defined as a function of ZC bonds or forward rates involved over the period[3] $[T_m, T_n]$:

$$S(t) = \frac{P(t,T_m) - P(t,T_n)}{B^S(t)} = \sum_{k=m}^{n-1} \alpha_k(t) F_k(t),$$

where the accrual period length is $\Delta T_i = T_{i+1} - T_i$, $B^S(t) = \sum_{i=m+1}^{n} \Delta T_{i-1} P(t,T_i)$ is the annuity of the swap and the stochastic weights are defined by $\alpha_k(t) = \frac{\Delta T_k P(t,T_{k+1})}{B^S(t)}$.

We ignore minor valuation adjustments needed for swaps which incorporate one or two business day payment delays to accommodate for the final rate fixing being unknown until the morning after for reformed benchmark rates.

Both forward and swap rates, as functions of ZC bonds, are quantities directly observable on financial markets. In its primary version, the LMM assumes a lognormal type dynamic for those rates, allowing to use the Black formula for the pricing of derivatives on forward rates (floorlet, caplet). In present DD-SV-LMM, the dynamics of forward rates are more complex. Still, as a function of forward rates, the dynamics of the swap rate—under an appropriate measure—can now be deduced from that of forward rates. For simplicity, we will only present the diffusion defining the swap rates. Note that the function linking swap rates and forward rates is independent of the modelling framework.

To price swap rates derivatives, it is convenient to work under a probability measure, making the swap rate a martingale. This is the case of the swap-forward measure, denoted by $\mathbb{Q}^S$, associated with the *numéraire* which appears in the definition of the swap rate. Calculation steps leading to the following dynamics are not elaborated here. The interested reader may refer to [WZ06] for a rigorous proof.

Let $\boldsymbol{Z}_t^S$ be a multidimensional Brownian motion under $\mathbb{Q}^S$ and W_t^S be a scalar Brownian motions under $\mathbb{Q}^S$. In the DD-SV-LMM, the evolution of the swap rate is specified by:

$$\begin{cases} \mathrm{d}S(t) = \sqrt{V(t)} \sum_{j=m}^{n-1} \omega_j(t)\big(F_j(t) + \delta\big)\, \gamma_j(t) \cdot \mathrm{d}\boldsymbol{Z}_t^S, \\ \mathrm{d}V(t) = \kappa(\theta - \xi^S(t)V(t))\mathrm{d}t + \epsilon\sqrt{V(t)}\mathrm{d}W_t^S, \end{cases} \quad (1)$$

where $V(t)$ is the instantaneous volatility of the swap rate parametrised by κ, θ, ϵ that are positive parameters, and γ_j is a bi-dimensional deterministic function modelling a part of the volatility structure of the j^{th} forward rate F_j:

$$\omega_j(t) = \alpha_j(t) + \frac{\Delta T_j}{1 + \Delta T_j F_j(t)} \sum_{p=m}^{j-1} \alpha_p(t) \left(F_p(t) - S_{m,n}(t)\right),$$

$$\xi^S(t) = 1 + \frac{\epsilon}{\kappa} \sum_{j=m}^{n-1} \alpha_j(t) \sum_{k=1}^{j} \frac{\Delta T_k (F_k(t) + \delta)}{1 + \Delta T_k F_k(t)} \rho_k(t) \|\gamma_k(t)\|,$$

ρ_k is a deterministic function accounting for the correlation between k^{th} forward rate F_k and stochastic volatility $V(t)$, and $\delta \geq 0$ is the shift coefficient allowing us to generate negative rates.

[3] Note that this formula is not valid in the case of mismatched discounting.

Note that, for numerical experiments, we have worked with a piecewise constant deterministic volatility structure set as:

$$\gamma_j(t) = \beta_j g(T_j - T_i) \text{ for } T_i \le t < T_{i+1},$$

where β_j is a bi-dimensional vector accounting for the inter-forward correlation structure, $g(u) = (a + bu)e^{-cu} + d$ for any $u \ge 0$, and:

$$\rho_j(t)\|\gamma_j(t)\| = \frac{\rho}{\sqrt{2}}\left(\gamma_j^{(1)}(T_j - T_i) + \gamma_j^{(2)}(T_j - T_i)\right).$$

In the end, the parameters of the model to be calibrated are $(a, b, c, d, \kappa, \theta, \epsilon, \rho)$. δ has been fixed in our experiments.

As it stands, the model for the swap rate is too complex to be usable. A common practice is to approximate the diffusion (1) by replacing the stochastic quantities appearing in the definition of ω_j and ξ^S above by their initial values by assuming that they are of low variability. A recent study on the validity of this assumption has been made in [MEH21]. It amounts to approximate swap rates as linear combinations of forward rates with stochastic volatility whose drift is deterministic. The final model on the swap rates with which we will work in the remainder of this paper is the so-called “normal” frozen dynamics:

$$\begin{cases} \mathrm{d}S(t) = \sqrt{V(t)}\sum_{j=m}^{n-1}\omega_j(0)(F_j(0)+\delta)\gamma_j(t)\cdot \mathrm{d}\boldsymbol{Z}_t^S, \\ \mathrm{d}V(t) = \kappa(\theta - \xi_0^S(t)V(t))\mathrm{d}t + \epsilon\sqrt{V(t)}\mathrm{d}W_t^S, \end{cases} \quad (2)$$

where:

$$\xi_0^S(t) = 1 + \frac{\epsilon}{\kappa}\sum_{j=m}^{n-1}\alpha_j(0)\sum_{k=1}^{j}\frac{\Delta T_k(F_k(0)+\delta)}{1+\Delta T_k F_k(0)}\rho_k(t)\|\gamma_k(t)\|.$$

For the sake of completeness, let us mention that, in some specification, Equation (2) may represent the dynamics of the logarithm of the swap rate (and not the swap rate itself); we would refer to this modelling framework as “lognormal” freezing.

In the following, we will denote by ϕ the moment-generating function of $S(t)$, defined by $\phi_{m,n}(x) = \mathbb{E}_t^S\left[e^{xS(T_m)}\right]$. When the swap rate is described by Equation (2), the characteristic function can be analytically derived using the affine property of the model (see [DABB17] for more details).

The knowledge of this characteristic function is useful for pricing methods. Indeed, the spot price $PS(0; T_m, T_n, K)$ of a payer’s swaption of maturity T_m, tenor $T_n - T_m$ and strike K, is given by:

$$\begin{aligned} PS(0; T_m, T_n, K) &= B^s(0)\mathbb{E}^S[\max(S(T_m) - K, 0)] \\ &= B^s(0)\left(\mathbb{E}^S\left[S(T_m)\mathbb{1}_{\{S(T_m)\ge K\}}\right] \right. \\ &\quad \left. - K\mathbb{E}^S\left[\mathbb{1}_{\{S(T_m)\ge K\}}\right]\right). \end{aligned}$$

It can be proved (see [WZ06] and references therein for the details of the calculations) that each expectation in the above equation can be expressed as a complex integral of the moment-generating function ϕ extended over a domain of the complex field $\mathcal{C} \subset \mathbb{C}$. Namely:

$$\mathbb{E}^S\left[S(T_m)\mathbb{1}_{\{S(T_m)\ge K\}}\right] = (S(0)+\delta)\left(\frac{1}{2} + \frac{1}{\pi}\int_0^\infty \Re\mathrm{e}\left(\frac{e^{-iu(K+\delta)}\phi(u-i)}{iu}\right)\mathrm{d}u\right),$$

$$\mathbb{E}^S\left[\mathbb{1}_{\{S(T_m)\ge K\}}\right] = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \Re\mathrm{e}\left(\frac{e^{-iu(K+\delta)}\phi(u)}{iu}\right)\mathrm{d}u,$$

where i denotes the imaginary unit defined as satisfying $i^2 = -1$, $\Re\mathrm{e}$ denotes the real part of a complex number for $z = z_1 + iz_2 \in \mathcal{C}$, $\Re\mathrm{e}(z) = z_1$ and $\mathbb{1}_A$ is the indicator function of the event A, defined by $\mathbb{1}_A(\omega) = 1$ if $\omega \in A$, and is equal to 0 otherwise.

Historically, this pricing method has been first used for pricing of the option in the popular Heston model widely used for equity modelling (see [HES93]). Therefore, in the numerical experiments presented in the final section of this paper, we will refer to it as the “Heston pricing method.”

Cosine expansion

The previous Heston-like formula for swaption prices is approximated in practice using Gaussian quadrature. This requires an important number (several hundreds) of calls to the integrated function, which is computed by solving Riccati equations. All in all, the Heston-like approach is quite time-consuming. Alternatively, the spot price of a swaption can be written as the integral of the payoff function of the option against the density function of the swap rate process:

$$\begin{aligned} PS(0; T_m, T_n, K) &= B^s(0)\mathbb{E}^S[\max(S(T_m) - K, 0)] \\ &= B^s(0)\int_{-\infty}^{+\infty}\max(s - K, 0)f_{T_m}(s)\mathrm{d}s \quad (3) \end{aligned}$$

where $f_{T_m}(s)$ is the density function of $S(T_m)$.

A straightforward calculation of this expectation is, however, impossible because there is no analytical expression for the swap rate density function in the DD-SV-LMM framework. One way to circumvent this issue is to derive an approximated density function using a Fourier expansion series.

The cos-pricing method relies on the fact that the characteristic function $\omega \in \mathbb{R} \mapsto \phi(i\omega)$ and the probability density function f_T are linked together by Fourier transformation through the relations:

$$\phi(i\omega) = \int_{-\infty}^{\infty} e^{ix\omega} f_T(x)\mathrm{d}x, \quad (4)$$

$$f_T(x) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-i\omega t}\phi(i\omega)\mathrm{d}\omega.$$

The Fourier transform of periodic functions can be expressed analytically as the sum of a convergent series. This property is used to derive alternative swaption prices. We first describe the Fourier transform of functions with finite support.

DENSITY CALCULATION

A function f is said to be supported on the interval $[a,b]$ if $f(x)=0$ for all $x \notin [a,b]$; it is said to be finitely supported over $[a,b]$ if $\max(|a|,|b|)<\infty$, which will be assumed in the following. Primarily used for periodic functions approximation with cosine series, the Fourier series can provide optimal approximation of finitely supported functions. A function f finitely supported over $[a,b]$ can be considered as a $(b-a)$-periodic function. A Fourier series expansion can then be performed on f. For functions supported on $[0,\pi]$, the cosine expansion reads:

$$f(\theta)=\frac{A_0}{2}+\sum_{k=0}^{\infty} A_k\cos(k\theta),$$

$$\text{with } A_k=\frac{2}{\pi}\int_0^{\pi} f(\theta)\cos(k\theta)\mathrm{d}\theta.$$

For a function f supported on arbitrary compact interval $[a,b]$, the expression of the cosine expansion can be deduced using that $[a,b]$ as a linear transformation of interval $[0,\pi]$:

$$f(x)=\frac{A_0}{2}+\sum_{k=1}^{\infty} A_k\cos\left(k\pi\frac{x-a}{b-a}\right),$$

$$\text{where } A_k=\frac{2}{b-a}\int_a^b f(x)\cos\left(k\pi\frac{x-a}{b-a}\right)\mathrm{d}x.$$

In practice, these expansions are truncated to a finite summation (or expansion) of order $N\in\mathbb{N}$, so that numerically it is computed:

$$f(\theta)\approx\frac{A_0}{2}+\sum_{k=1}^{N} A_k\cos\left(k\pi\frac{x-a}{b-a}\right). (5)$$

NUMERICAL ILLUSTRATION

We illustrate the cosine approximation series on normal and uniform probability distributions. Observe that the uniform distribution is indeed a finitely supported density over [0,1] while the normal one is not. Density approximations are plotted at different truncation orders N (see formula [4]) to demonstrate the ability of accurately approximate target densities.

FIGURE 1: APPROXIMATION OF A NORMAL DISTRIBUTION WITH COSINE SERIES AT DIFFERENT ORDERS

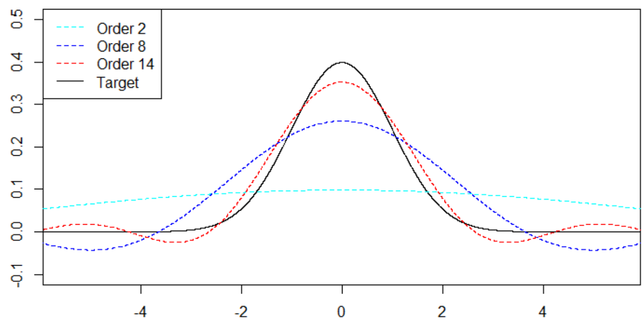

It is seen in Figure 1 that a cosine-series efficiently approximates the normal distribution density function, as the successive approximations converge to the target density (plotted in black).

FIGURE 2: APPROXIMATION OF A UNIFORM DISTRIBUTION WITH COSINE SERIES AT DIFFERENT ORDERS

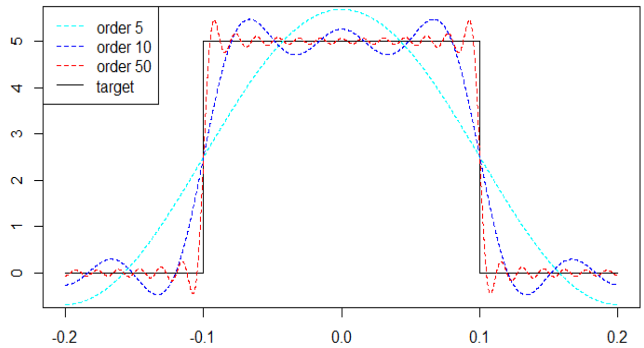

In Figure 2, the approximation of the uniform density is illustrated. We observe that there still is a convergence towards the target density (black line) but that the convergence is much slower due to the discontinuity of the uniform density function. Oscillations appear around the discontinuity points (also known as the Gibbs phenomenon).

COSINE EXPANSION FOR SWAPTION PRICING

For our problem of swaption pricing, we aim at applying such cosine expansions to compute the swap rate density function f_{T_m}. However, there is no reason for f_{T_m} to be supported on finite intervals. To apply the cosine expansions to the swap rate density, we assume that we can approximate f_{T_m} by a function that is supported on a finite compact interval $[a,b]$, denoted by $f_{T_m}^{(a,b)}$. To apply this technique, $a,b\in\mathbb{R}$ must be chosen such that the swap rate density is negligible on both $]-\infty,a]$ and $[b,+\infty[$, which in turns requires that $\lim_{x\to+\infty} f(x)=\lim_{x\to-\infty} f(x)=0$. These assumptions would allow us to:

- Ensure the target density function can be approximated by a finitely supported function without losing too much information on the swap rate distribution
- Ensure the Fourier transform of the approximating function (that is, the approximating characteristic function) is close to the Fourier transform of the original density function (that is, the target characteristic function)

By properly chosen $a,b\in\mathbb{R}$, f_{T_m} can be approximated by $f_{T_m}^{(a,b)}$, whose cosine series is characterised by the coefficients:

$$A_k=\frac{2}{b-a}\Re\left\{\left(\int_a^b e^{ik\pi\frac{x}{b-a}} f_{T_m}^{(a,b)}(x)\mathrm{d}x\right)\cdot e^{-ik\pi\frac{a}{b-a}}\right\} (6).$$

We will refer to this approximation method as the "cosine approximation series" in this paper. Furthermore, the characteristic function of the model can be recovered using Equation (4) on a finite number of points. For all $k \in \mathbb{N}$:

$$\phi\left(\frac{ik\pi}{b-a}\right) = \int_{-\infty}^{\infty} e^{ik\pi\frac{x}{b-a}} f_T(x)dx \approx \int_a^b e^{ik\pi\frac{x}{b-a}} f_T(x)\mathrm{d}x,$$

which can be injected in Equation (6) to approximate the Fourier coefficients as:

$$A_k = \frac{2}{b-a}\Re\mathrm{e}\left\{\phi\left(\frac{ik\pi}{b-a}\right)\cdot e^{-ik\pi\frac{a}{b-a}}\right\}.$$

As with Equation (5), these approximating coefficients yield a relatively simple approximation of the target density function.

Let us mention that the cosine expansion can also be applied in a lognormal framework. In this version of the DD-SV-LMM, the characteristic function of $\ln(S(T))$ is known instead of that of $S(T)$. This requires a minor adjustment of the model defined by the dynamics in Equation (2). Similar calculations lead to the following expression for a swaption price in a lognormal framework:

$$\mathbb{E}^S\left[\left(e^{\frac{S(T)}{K}}-1\right)^+\right] = \frac{e^b - b - 1}{(b-a)} + \sum_{k=1}^N A_k V_k,$$

$$\text{with } A_k = \frac{2}{b-a}\Re\mathrm{e}\left\{\phi\left(ik\pi\frac{x}{b-a}\right)\cdot e^{-ik\pi\frac{a}{b-a}}\right\},$$

$$V_k = \frac{1}{1+\left(\frac{b-a}{k\pi}\right)^2}\left[(-1)^k e^b - \cos\left(-k\pi\frac{a}{b-a}\right) + \frac{b-a}{k\pi}\sin\left(-k\pi\frac{a}{b-a}\right)\right].$$

Whatever the selected modelling framework is, the cosine expansion for swaption pricing can be summarised as follows:

1. Computation of the moment-generating function of the model on a finite set of points (crucial step for reduction of computational time).
2. Computation of the cosine series coefficients.
3. Computation of the density function.

SWAPTION PRICING

We now present how the cosine approximation series can be employed in our context of swaption pricing. The method consists in approximating the swap rate density function using Equation (5) and injecting it into the swaption price formula of Equation (3).

We then obtain an approximated swaption price as

$$\mathbb{E}^S[\max(S(T)-K,0)] \approx \int_a^b \max(x-K,0)\left(\frac{1}{b-a} + \sum_{k=1}^N A_k \cos\left(k\frac{x-a}{b-a}\right)\right)\mathrm{d}x.$$

where $a < K < b$.

The previous expressions can be further developed to get an analytical formula for the approximated swaption price:

$$\mathbb{E}^S[\max(S(T)-K,0)] = \frac{(b-K)^2}{2(b-a)} + \sum_{k=1}^N A_k V_k,$$

$$\text{with } A_k = \frac{2}{b-a}\Re\mathrm{e}\left\{\phi\left(ik\pi\frac{x}{b-a}\right)\cdot e^{-ik\pi\frac{a}{b-a}}\right\},$$

and $V_k = \int_a^b \max(x-K,0)\, A_k \cos\left(k\frac{x-a}{b-a}\right)\mathrm{d}x = \left(\frac{b-a}{k\pi}\right)^2\left((-1)^k - \cos\left(k\pi\frac{K-a}{b-a}\right)\right)$, for $k \leq N$.

Swaptions can now be priced with N calls to the characteristic function. The optimal series truncation order strongly depends on the density function truncation interval $[a,b]$. However, there are no analytical results backing the appropriate choice for a and b. This is an important issue because the choices of a and b have crucial implications on the convergence of the method as they can impair both convergence speed and accuracy, as exposed in the following figures. [FAN10] proposed an interval based on the cumulants of the distribution:

$$[a,b] \triangleq \left[\xi_1 - L\sqrt{\xi_2+\sqrt{\xi_4}}, \xi_1 + L\sqrt{\xi_2+\sqrt{\xi_4}}\right] \quad \text{with } L = 10, (7)$$

$$\xi_1 = \mathbb{E}^S[S(T)] = S(0),$$

$$\xi_2 = \mathbb{E}^S[(S(T)-\mathbb{E}^S[S(T)])^2],$$

$$\xi_4 = \mathbb{E}^S[(S(T)-\mathbb{E}^S[S(T)])^4] - 3(\xi_2)^2.$$

This choice has proven to perform in a satisfying way in multiple numerical applications.

ILLUSTRATIVE APPLICATIONS

To illustrate the impact of the choice of those coefficients a, b, we depict in Figure 3 a cosine expansion performed on a given centred Gaussian distribution of standard deviation $\sigma = 1$ with a fixed truncation order N. Three expansions have been performed, with different values of a and b: namely, we consider $(a,b) \in \{(3\sigma,-3\sigma),(5\sigma,-5\sigma),(10\sigma,-10\sigma)\}$ and we compare them to the exact density functions of the Gaussian distribution (analytically known).

FIGURE 3: APPROXIMATION OF A NORMAL DISTRIBUTION WITH COSINE SERIES AT ORDER 8 FOR DIFFERENT TRUNCATION INTERVALS

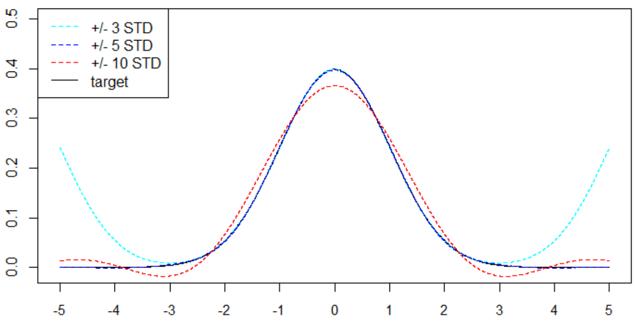

We observe that a tight interval $[a, b]$ leads to loss of information on the tails of the distribution while a loose interval induces a slower convergence towards the reference density.

In an experiment illustrated in Figure 3, the distribution was analytically known and thus setting a convenient interval $[a, b]$ in the function of the parameter of the target density was possible. In practice, this is not possible. In Figure 4, we provide an experiment in which the cosine expansion is applied to the swaption pricing problem. Namely, using parameters previously calibrated with the reference Heston-like method, 258 swaptions are priced with the $[-1, 1]$ interval on one hand and with the adjusted interval based on the cumulants defined above in Equation (7) on the other hand. The sum of the squared error between the prices computed with the reference method and the prices computed with the cos-pricing method is represented with respect to the expansion order N in Figure 4. The parameters used to perform this numerical experiment are the following:

$$(a, b, c, d, \kappa, \theta, \epsilon, \rho, \delta) = (0.080, 0.00022, 0.053, 0.025, 0.10, 0.46, 0.31, 0.99, 0.1).$$

FIGURE 4: ACCURACY OF THE COS-PRICING METHOD FOR DIFFERENT APPROXIMATION INTERVALS WITH RESPECT TO THE EXPANSION ORDER

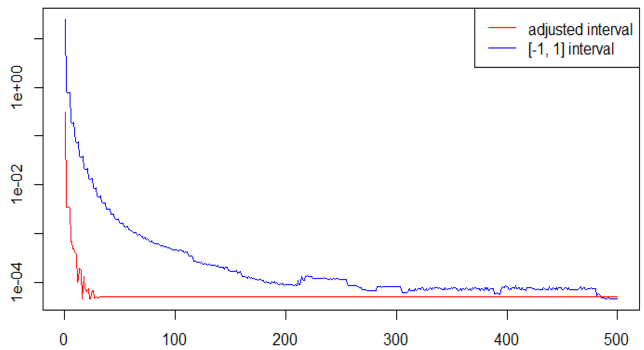

Both intervals provide relatively fast convergence towards the reference prices; yet the adjusted interval built using cumulants provides much faster convergence than the loose interval $[-1,1]$. When handling the cos-expansion series, the choice of the approximation interval must be systematically monitored because it is a key leverage to improve the method performances.

We then perform simple experiments in which we compute swaptions prices by assuming the swap rate is distributed following a Gaussian (with mean 0 and variance $1/300$) or uniform (over $[-0.1, 0.1]$) densities. In these cases, the value of $\mathbb{E}^S[(S(T) - K)^+]$ can be computed via closed-form formulas. Those exact prices are then compared to what is computed using the cosine expansion approximation in Equation (5), at different expansion orders (i.e., different values of N).

FIGURE 5: COS-PRICING WHEN THE SWAP RATE IS DISTRIBUTED ACCORDING TO A GAUSSIAN DENSITY

Strike	Order 2	Order 4	Order 8	Target
-0.02	0.03506	0.03537	0.03441	0.03440
0.00	0.02415	0.02415	0.02305	0.02303
0.02	0.01506	0.01537	0.01441	0.01440

FIGURE 6: COS-PRICING WHEN THE SWAP RATE IS DISTRIBUTED ACCORDING TO A UNIFORM DENSITY

Strike	Order 5	Order 10	Order 50	Target
-0.02	0.03533	0.03599	0.03600	0.03600
0.00	0.02420	0.02494	0.02500	0.02500
0.02	0.01533	0.01599	0.01600	0.01600

It is first observed in the tables in Figures 5 and 6 that in all experiments, the cosine expansion method does converge towards the exact prices. Moreover, this method provides very accurate results as of the very first expansion orders.

DD-SV-LMM calibration

Before going any further, let us specify our terminology. In the rest of this paper: prices induced by pricing formulas derived under the DD-SV-LMM framework (such as the cos-pricing or the Heston-like methods presented in this paper) will be referred to as “model prices.” Prices extracted from the market will be referred to as “market prices.” Finally, prices obtained by Monte Carlo simulations will be referred to as “Monte Carlo prices.”

CALIBRATION FRAMEWORK

As stated in the Motivations section above, the DD-SV-LMM framework is calibrated on swaption prices quoted on the market. In fact, an adequation metric between market and model prices is defined and minimised. Because swaptions are quoted through implied volatilities it would seem natural to set volatilities as reference quantities to replicate. Yet prices have been chosen instead of volatility because it is the raw output of pricing formulas.

The metric used for calibration is the root mean square relative error between model and market prices defined as follows:

$$\text{RMSE} = \sqrt{\frac{1}{N_{\text{data}}} \sum \left(\frac{\text{Price}_{\text{model}} - \text{Price}_{\text{market}}}{\text{Price}_{\text{market}}} \right)^2}$$

Several calibration experiments have been undertaken on multiple sets of market prices. Those calibration data gather three economies considered at five dates, including the turmoiled data of the first quarter (Q1) of 2020 with the COVID-19 pandemic outbreak and Q1 2022 with the trigger of the war in Ukraine. At each date, the number of euro market data that are replicated is of $N_{\text{data}} = 258$, composed of at-the-money normal implied volatilities whose maturities and tenors range within $\{1,2,3,4,5,7,10,15,20,25,30\}$ and away-from-the-money data of tenor 10 years and (relative) strikes within $\{\pm 25 \text{ bips}, \pm 50 \text{ bips}, \pm 100 \text{ bips}, \pm 150 \text{ bips}, \pm 200 \text{ bips}\}$.

CALIBRATION RESULTS

The first concrete improvement of the cosine expansion method is the calibration speed. Indeed, pricing a set of swaption took 4.56 milliseconds and calibration took five seconds, with a 30th-order cosine-expansion—as observed in Figure 4 above, the expansion error settles down as of $N = 30$, motivating our choice—whereas the reference method took, respectively, 17.2 milliseconds (durations obtained by averaging over $2{,}56 \cdot 10^6$ calls to the pricing function) and 16 seconds to achieve the same tasks.

One can worry that this gain in computational time is achieved at the cost of a lower quality in the fit of market data. However, this is not the case and the cosine expansion method performs very well also in terms of data replication. Indeed, cosine expansion series allow for an accurate fit of market prices with an average root mean square error (RMSE) over 15 calibrations of 1.50% for cosine approximation in a lognormal framework that should be compared with the RMSE of 1.49% obtained with the reference method. Note that those 15 calibration experiments have been realised using different market data (at different dates and on different economies) so that we can assess the stability of the proposed method. We provide in Figure 7 the box plots of the calibration errors obtained as outputs of these 15 minimisation procedures we have considered. Namely, the box plots contain RMSE between market data and model quantities obtained as outputs of calibrations (on the left of Figure 7) and the RMSE between model quantities and simulated quantities (on the right of Figure 7). It includes the fit accuracy for the reference Heston-like method and the cosine expansion method applied in lognormal framework.

FIGURE 7: ADEQUATIONS METRICS BETWEEN MARKET/MODEL AND MODEL/MONTE CARLO VALUES

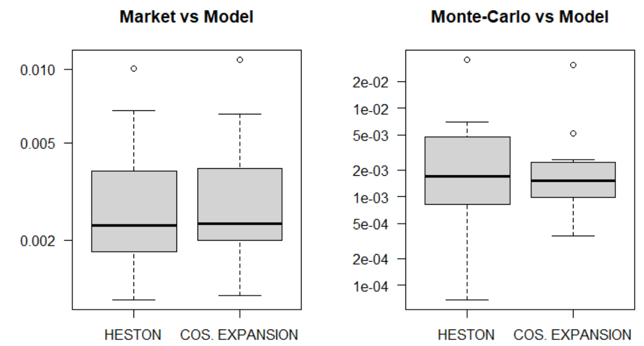

One can see on the plot that the two versions of the cos-pricing perform well enough to compare to the reference Heston-like method. The replication of market data by the model (observed in the plot "Market vs. Model") is very close to what was obtained with the reference method. Moreover, cosine expansion seems to induce fewer variances on the distance between model prices and volatilities and Monte Carlo (simulated) ones. In the plots in Figures 8 and 9, we provide the replication (by calibrated models) of prices and implied volatilities of swaptions observed in the euro market, as of 31 March 2022, for 1 year (1Y) x 10 years (10Y) (Figure 8) and for 10 years x 10 years (with the convention, maturity x tenor, in Figure 9).

FIGURE 8: PRICE AND VOLATILITY SMILE FOR A 1Y X 10Y SWAPTION

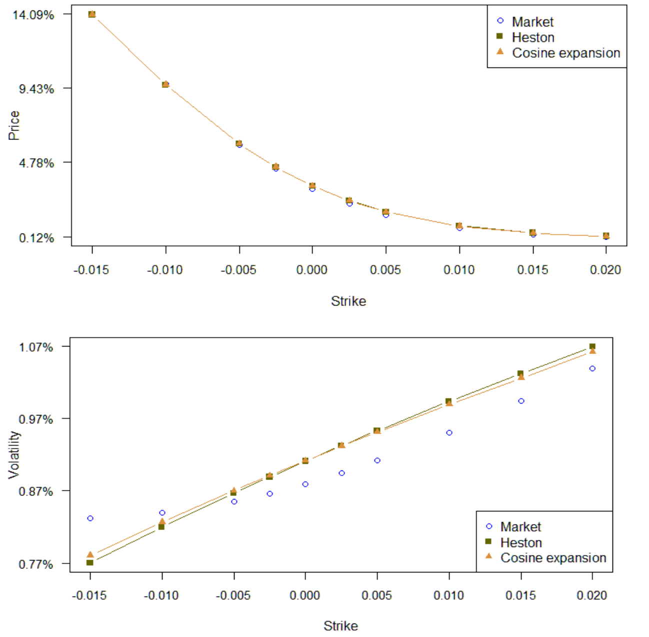

FIGURE 9: PRICE AND VOLATILITY SMILE FOR A 10Y X 10Y SWAPTION

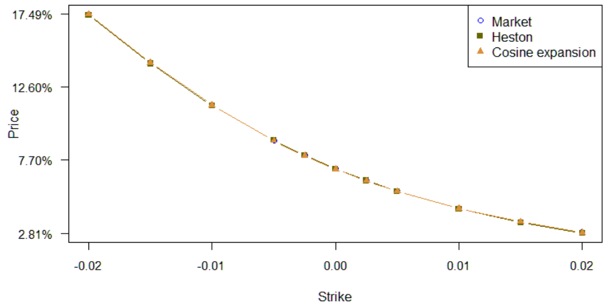

FIGURE 9: PRICE AND VOLATILITY SMILE FOR A 10Y X 10Y SWAPTION (CONTINUED)

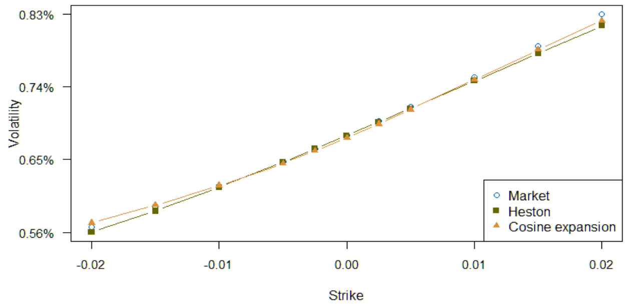

All the experienced methods provide satisfactory results in terms of data replication. The nonlinear transformation between prices and implied volatilities preserves the fact that the lognormal framework outperforms the normal one.

Conclusion

The development of sophisticated interest rates models improved the ability to replicate market data at the expense of a significant loss of analytical tractability and an increase of the computational requirements. The fundamental need of economic scenario generators in the insurance business pushes towards a constant improvement in calibration performance. With that in mind, this paper provides an alternative pricing method based on density cosine-series approximations applied to swaptions under the DD-SV-LMM dynamics. This method offers a significant gain in computation time efficiency. Its accuracy has not been directly verified but market price consistency has been maintained through calibration. The work presented above can be extended to any financial model dedicated to other risk factors and different option types, at the sole conditions that the moment-generating function of the underlying asset is available.

REFERENCES

1. [DABB17] Devineau, Arrouy, Bonnefoy and Boumezoued, *Fast calibration of the Libor Market Model with Stochastic Volatility and Displaced Diffusion*, Journal of Industrial and Management Optimization, 2020.
2. [WZ06] Wu and Zhang, *Libor Market Model with Stochastic Volatility*, Journal of Industrial and Management Optimization, 2006.
3. [FAN10] Fang, *The COS Method: An Efficient Fourier Method for Pricing Financial Derivatives*, 2010.
4. [HES93] Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*. The Review of Financial Studies, 1993, vol. 6, no 2, p. 327-343.
5. [MEH21] Sophian Mehalla, *Interest rates for insurance: approximations and calibrations of models*, 2021.

CONTACTS

Sophian Mehalla
sophian.mehalla@milliman.com

Lucien Morice
lucien.morice@milliman.com