numerix

Numerix **Journal**

Vol. 2 No. 1

May 2015

The Numerix Journal

Numerix Journal
Wednesday 27th May, 2015

Table of Contents

The Free Boundary SABR: Natural Extension to Negative Rates

Alexandre Antonov* Michael Konikov† Michael Spector‡

Abstract

In the current low-interest-rate environment, extending option models to negative rates has become an important issue. This paper describes one such extension of the widely used SABR model. *We stress that our solution is more natural and attractive than the shifted SABR model.* An exact formula is derived for option prices in the case of zero correlation between the rate and its volatility. For nonzero correlation, a mapping procedure onto a mimicking zero-correlation model is applied. Analytical results from the suggested free-boundary SABR model are compared with Monte Carlo simulation results.

1 Introduction

The SABR process with parameters $(F_0, v_0, \beta, \rho, \gamma)$[1] [7] for a rate F_t and its volatility v_t has the SDE

$$dF_t = F_t^\beta \, v_t \, dW_1, \tag{1.1}$$

$$dv_t = \gamma \, v_t \, dW_2, \tag{1.2}$$

with correlation $\mathbb{E}[dW_2 \, dW_1] = \rho dt$ and power $0 \le \beta < 1$. The solution is not uniquely defined by the SDE—we also need to impose boundary conditions. The standard choice is to assume that the boundary at zero is absorbing, which enforces positivity and martingality of the rate. See [2], [4], [6], [8], [9], [10], [13] for further references.

The SABR model is primarily used for volatility cube interpolation and for pricing CMS products by replication with vanilla options. It is also used in term-structure models, e.g., [12], [14].

When the SABR model was first introduced, rate positivity seemed like a reasonable and attractive property. In the current market environment, where rates are extremely low and sometimes even negative, it is important to extend the SABR model to negative rates. For example, Figure 1 shows the historical evolution of Swiss Franc (CHF) interest rates. One observation is that rates reached as low as -2%. Another important observation is that the rates "stick" to zero for certain periods of time, suggesting that their probability density functions have a singularity at zero.

The simplest way to take negative rates into account is to shift the SABR process

$$dF_t = (F_t + s)^\beta \, v_t \, dW_1,$$

where s is a deterministic positive shift. This moves the lower bound of F_t from 0 to $-s$.

One can either include the shift in the calibration parameters $(v_0, \beta, \rho, \gamma, \mathbf{s})$ or fix it prior to calibration (e.g., to 2% in the case of Swiss Franc short rates). Each alternative has drawbacks.

Calibrating the shift does not really introduce a new degree of freedom: its influence on the skew is very similar to the power β and may result in an identification problem. Moreover, one still needs to provide its lower bound (say 2%) as an input to a numerical solver. Otherwise, we risk reducing the rate span. This means that we still end up with an unquoted input to our model.

*Numerix, SVP of Quantitative Research, antonov@numerix.com
†Numerix, Executive Director of Quantitative Development, mkonikov@numerix.com
‡Numerix, Director of Quantitative Research, mspector@numerix.com

[1] Sometimes α is used instead of v_0.

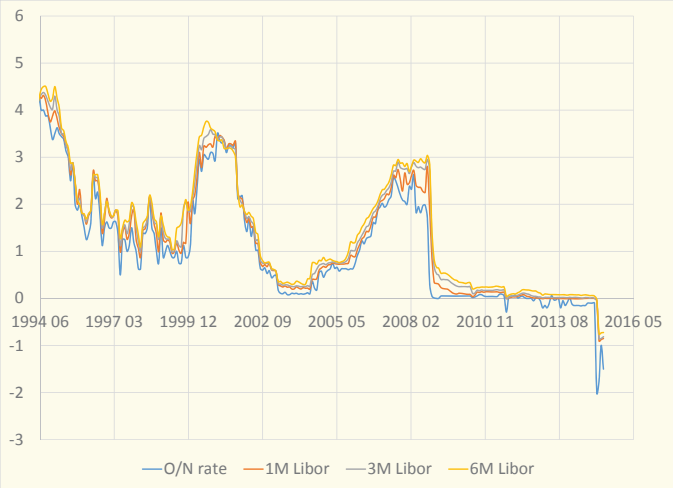

Figure 1: Swiss Franc interest rates.

If we fix the shift from the very beginning and calibrate the standard parameters $(v_0, \beta, \rho, \gamma)$, there are still drawbacks. There is always the danger that rates can go lower than anticipated, in which case we would need to change this parameter accordingly. This can result in a jump in the other SABR parameters as a calibration response to such readjustment. As a consequence, we can get jumps in the values/Greeks of the trades dependent on swaption or cap volatilities. To cover for potential losses in such situations, traders are likely to be asked to reserve part of their P&L. Also, having the swaption prices bounded from above (due to the rate being bounded from below) can lead to situations when the shifted SABR model cannot attain market prices. In sum, we need a more natural and elegant solution for permitting negative rates.

For $\beta = 0$, the normal SABR model, $dF_t = v_t\, dW_1$, allows the rates to become negative when a free boundary condition is enforced. Below, we come up with a generalization of this model

$$dF_t = |F_t|^\beta\, v_t\, dW_1$$

with $0 \le \beta < \frac{1}{2}$ and a *free* boundary. As we will see, such a model allows for rates that can be negative and exhibit a certain "stickiness" at zero.

In what follows, we consider only the $F_0 > 0$ case (unless explicitly stated otherwise). When $F_0 < 0$, we note that $\tilde{F}_t = -F_t$ satisfies the SABR SDE with parameters $(-F_0, v_0, \beta, -\rho, \gamma)$, and the time value of a European option on F_t struck at K equals that of an option on $\tilde{F}_t$ struck at $-K$.

To get intuition about the free boundary, we start with a CEV example $dF_t = |F_t|^\beta\, dW$ and study the PDF and option prices. Then we switch to the SABR model with a free boundary condition and present an exact solution for the zero-correlation case. For the general case, we show an accurate approximation for European options prices. We demonstrate that the exact formula as well as its approximation can be presented in terms of a 1D integral over elementary functions, making it well-suited for fast calibration.[2] We finish with simulation schemes and numerical results.

2 CEV Process

To aid with intuition, we consider the CEV model $dF_t = F_t^\beta\, dW$ with $0 \le \beta < 1$. The forward Kolmogorov (FK) equation on the density $p(t, f)$,

$$p_t - \frac{1}{2}\left(f^{2\beta}\, p\right)_{ff} = 0,$$

[2]Note that the SABR approximation [7] based on the heat-kernel expansion cannot be applied to the free SABR because it does *not* tale into account the boundary conditions.

has two types of solutions, depending on the boundary conditions; fixing the PDE (or SDE) alone is not sufficient to uniquely define the solution. One can show (e.g., [5]) that there are two distinct solutions with asymptotics $p_A \sim f^{1-2\beta}$ and $p_R \sim f^{-2\beta}$. We call the first solution *absorbing* and the second one *reflecting*. The latter exists only for $\beta < \frac{1}{2}$; otherwise, the norm around zero diverges.

The asymptotics are closely related to conservation laws, which can be obtained by integrating the FK equation by parts against some payoffs $h(f)$. Consider first the norm case $h(f) = 1$. It is easy to see the asymptotic behavior of the absorbing solution leads to nonconservation of the norm, while the reflecting solution conserves the norm. For the first moment conservation, we take $h(f) = f$ and deduce that the asymptotics of the reflecting solution leads to nonconservation of the first moment (i.e., nonmartingality), while the absorbing solution is a martingale.

The explicit PDFs for the CEV process are known (see [11] and [5]) in terms of the modified Bessel functions, which permits us to calculate a call option time-value via the time integral without the boundary term:

$$\mathcal{O}(T,K) = \mathbb{E}[(F_T - K)^+] - (F_0 - K)^+ = \frac{1}{2}K^{2\beta}\int_0^T dt\, p(t,K). \tag{2.1}$$

As explained in [5], this is not the case for put options, where a boundary term is present.

Below we will need option prices for absorbing/reflecting solutions as 1D integrals (see [3] and [5]). These are given by

$$\mathcal{O}_{A/R}(T,K) = \frac{\sqrt{KF_0}}{\pi}\left(\int_0^\pi \frac{\sin(|\nu|\theta)\sin(\theta)}{b-\cos(\theta)} e^{-\frac{\bar{q}(b-\cos(\theta))}{T}} d\theta \right.$$
$$\left. + \sin(|\nu|\pi)\int_0^\infty \frac{e^{\mp|\nu|x}\sinh(x)}{b+\cosh(x)} e^{-\frac{\bar{q}(b+\cosh(x))}{T}} dx\right) \tag{2.2}$$

for an index $\nu = -\frac{1}{2(1-\beta)}$ and parameters

$$\bar{q} = q_0\, q_K, \quad b = \frac{q_0^2 + q_K^2}{2\, q_0\, q_K}, \quad q_0 = \frac{F_0^{1-\beta}}{1-\beta} \quad \text{and} \quad q = \frac{K^{1-\beta}}{1-\beta}.$$

Now consider an extension of the CEV model to the entire real line by modifying the SDE to

$$dF_t = |F_t|^\beta\, dW \tag{2.3}$$

for $0 \le \beta < \frac{1}{2}$. The corresponding FK equation is

$$\partial_t\, p(t,f) = \frac{1}{2}\left(|f|^{2\beta}\, p(t,f)\right)_{ff}. \tag{2.4}$$

A *norm conserving* and *martingale* solution that satisfies the FK equation with the initial condition $p(0,f) = \delta(f - F_0)$ can be constructed from the reflecting and absorbing solutions as

$$p(t,f) = \frac{1}{2}\left(p_R(t,|f|) + \text{sign}(f)\; p_A(t,|f|)\right). \tag{2.5}$$

We can get the same expression for the density with a purely probabilistic argument. For any $f > 0$, a reflecting path ending at f is equivalent to a free-boundary path ending at f or $-f$, and vice versa. For the absorbing case, we apply the reflection principal as usual, which involves taking the probability of a path ending at f and subtracting the probability of the path ending at f and touching zero, which is equivalent to a path ending at $-f$. Hence, we can write the linear system

$$p_R(t,f) = p(t,f) + p(t,-f),$$
$$p_A(t,f) = p(t,f) - p(t,-f).$$

Solving this yields the expression for the free boundary density. The solutions for typical parameters are shown in Figure 2.

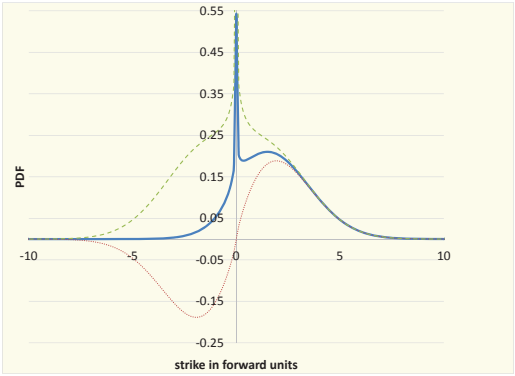

Figure 2: The blue solid line represents the free PDF, the red dotted line depicts the absorbing density expression $\text{sign}(f)\, p_A(t, |f|)$, while the green dashed line gives the symmetrical reflecting solution.

Note that the PDF diverges as $p(t, f) \sim |f|^{-2\,\beta}$ at zero. (The asymptote is inherited from the reflecting solution.) The observed singularity is quite natural: one can observe a "sticky" behavior of real rates near zero in, for example, the behavior of the CHF rate in Figure 1.

A call option payoff $h(f) = (f - K)^+$ leads to an option time value of

$$\begin{aligned}\mathcal{O}_F(T, K) &= \frac{1}{2}|K|^{2\beta} \int_0^T dt\, p(t, K) = \frac{1}{2}|K|^{2\beta} \int_0^T dt\, \frac{1}{2}\left(p_R(t, |K|) + \text{sign}(K)\; p_A(t, |K|)\right) \\ &= \frac{1}{2}\left(\mathcal{O}_R(T, |K|) + \text{sign}(K)\, \mathcal{O}_A(T, |K|)\right). \end{aligned} \tag{2.6}$$

Finally, we present the free CEV option integral. Its time value can easily be derived from the absorbing-reflecting solutions (2.2) and (2.6), yielding

$$\begin{aligned}\mathcal{O}_F(\tau, K) = \frac{\sqrt{|K F_0|}}{\pi}\Bigg(&\mathbf{1}_{K\geq 0} \int_0^\pi \frac{\sin(|\nu|\,\theta)\sin\theta}{b - \cos\theta}\, e^{-\frac{\bar{q}(b-\cos\theta)}{\tau}}\, d\theta \\ &+ \sin(|\nu|\,\pi) \int_0^\infty \frac{\left(\mathbf{1}_{K\geq 0}\cosh(|\nu|\,x) + \mathbf{1}_{K<0}\sinh(|\nu|\,x)\right)\sinh x}{b + \cosh x}\, e^{-\frac{\bar{q}(b+\cosh x)}{\tau}}\, dx\Bigg), \end{aligned} \tag{2.7}$$

where $\nu = -\frac{1}{2(1-\beta)}$ and

$$\bar{q} = \frac{|F_0\, K|^{1-\beta}}{(1-\beta)^2} \quad \text{with} \quad b = \frac{|F_0|^{2(1-\beta)} + |K|^{2(1-\beta)}}{2\,|F_0\, K|^{1-\beta}}.$$

We will use this formula to derive the analytics for the SABR model in the section below. Note that we put the absolute value of F_0 for symmetry with respect to the strike: F_0 is assumed to be positive, consistent with the remark in the introduction.

Regarding the sensitive region of small strikes and/or small rates, we see that the call option price (the full one, including the intrinsic value) is a smooth function of K and F_0 at zero. Thorough analysis reveals that the main terms of the expansion near zero are linear with the following terms of the order $|K|^{2(1-\beta)}$ for small strikes and $|F_0|^{2(1-\beta)}$ for small spots.

3 SABR

Now, let us come back to the SABR process (1.1)–(1.2). The standard choice of the absorbing boundary will be generalized to a *free* boundary. Namely, we will consider the SDE

$$dF_t = |F_t|^\beta\, v_t\, dW_1$$

for $0 \leq \beta < \frac{1}{2}$ (with the same process (1.2) for the stochastic volatility). Such a construction permits both *negative* rates and "stickiness" at zero.

Looking forward, we plot the SABR density function, which is shown in Figure 3 for the Input I parameters from Table 4. We also observe the singularity at 0, which reflects the "sticky" behavior of rates at zero. (See Figure 1.)

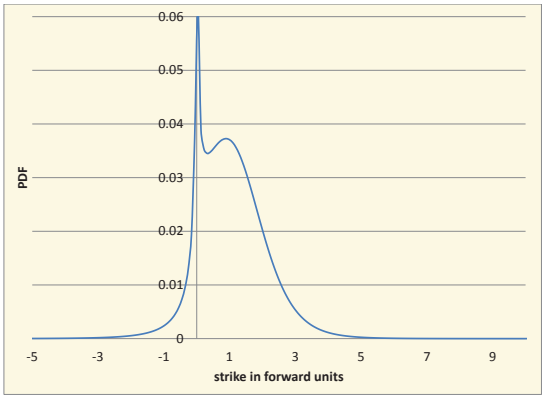

Figure 3: SABR model PDF for $T = 3Y$, $\beta = 0.25$.

3.1 Zero-Correlation Case

The zero-correlation free SABR model can be solved exactly. Indeed, the option price can be computed as

$$\mathcal{O}_F^{SABR}(T,K) = \mathbb{E}\left[\mathcal{O}_F^{CEV}(\tau_T,K)\right], \tag{3.1}$$

where $\mathcal{O}_F^{CEV}(\tau,K)$ is the free-boundary CEV option price (2.7), and the stochastic time $\tau_T = \int_0^T v_t^2 dt$ is the cumulative variance for the geometric Brownian motion v_t (1.2). The dependence on τ in both integrand terms of (2.7) is of the form $\exp(-\lambda/\tau)$. Thus, averaging over stochastic time, $\mathbb{E}\left[\mathcal{O}_F^{CEV}(\tau_T,K)\right]$, requires calculating the mean value $\mathbb{E}\left[\exp(-\lambda/\tau_T)\right]$.

The moment generating function (MGF) of the inverse stochastic time was derived in [3] and given by

$$\mathbb{E}\left[\exp\left(-\frac{\lambda}{\tau_T}\right)\right] = \frac{G(T\gamma^2,s)}{\cosh s} \quad \text{for} \quad s = \sinh^{-1}\left(\frac{\sqrt{2\lambda}\,\gamma}{v_0}\right).$$

The function

$$G(t,s) = 2\sqrt{2}\frac{e^{-\frac{t}{8}}}{t\sqrt{2\pi t}}\int_s^\infty du\, u\, e^{-\frac{u^2}{2t}}\sqrt{\cosh u - \cosh s}$$

has been introduced in [4]; it is closely related to the McKean heat kernel on the hyperbolic plane H^2. It is important to notice that, although the function $G(t,s)$ is a 1D integral, it can be *very* efficiently approximated by a closed formula (see [4]).

Thus, the exact option price for the zero correlation case can be presented as

$$\mathcal{O}_F^{SABR}(T,K) = \frac{1}{\pi}\sqrt{|KF_0|}\left\{\mathbf{1}_{K\geq 0}A_1 + \sin(|\nu|\pi)\,A_2\right\}$$

with the integrals

$$A_1 = \int_0^\pi d\phi \frac{\sin\phi \sin(|\nu|\phi)}{b-\cos\phi}\frac{G(T\gamma^2,s(\phi))}{\cosh s(\phi)}, \tag{3.2}$$

$$A_2 = \int_0^\infty d\psi \frac{\sinh\psi\left(\mathbf{1}_{K\geq 0}\cosh(|\nu|\psi) + \mathbf{1}_{K<0}\sinh(|\nu|\psi)\right)}{b+\cosh\psi}\frac{G(T\gamma^2,s(\psi))}{\cosh s(\psi)}. \tag{3.3}$$

Here s has the following parameterization with respect to ϕ and ψ:

$$\sinh s(\phi) = \gamma\, v_0^{-1}\, \sqrt{2\bar{q}(b - \cos\phi)},$$
$$\sinh s(\psi) = \gamma\, v_0^{-1}\, \sqrt{2\bar{q}(b + \cosh\psi)},$$

where $\bar{q}$ and b are the same as in the CEV free-boundary option.

3.2 General Correlation Case

As in [4], we approximate the general correlation option price by using the zero-correlation model

$$d\tilde{F}_t = |\tilde{F}|_t^{\tilde{\beta}}\, \tilde{v}_t\, d\tilde{W}_1,$$
$$d\tilde{v}_t = \tilde{\gamma}\, \tilde{v}_t\, d\tilde{W}_2,$$

with $\mathbb{E}[d\tilde{W}_1\, d\tilde{W}_2] = 0$. That is, we aim to find model parameters for $\tilde{F}$ so that

$$\mathbb{E}\left[(F_t - K)^+\right] \simeq \mathbb{E}\left[(\tilde{F}_t - K)^+\right].$$

For the *free* boundary, we reuse the same effective coefficients of the zero-correlation SABR model as in [4] for the *absorbing* boundary. The power and vol-of-vol are strike-independent with

$$\tilde{\beta} = \beta \quad \text{and} \quad \tilde{\gamma}^2 = \gamma^2 - \frac{3}{2}\left\{\gamma^2\rho^2 + v_0\gamma\rho\,(1-\beta)\,F_0^{\beta-1}\right\},$$

while the initial stochastic volatility is more complicated and *strike-dependent*. The $\tilde{v}_0$ parameters can be calculated as an expansion

$$\tilde{v}_0 = \tilde{v}_0^{(0)} + T\,\tilde{v}_0^{(1)} + \cdots. \tag{3.4}$$

The leading volatility term can be expressed as

$$\tilde{v}_0^{(0)} = \frac{2\,\Phi\,\delta\tilde{q}\,\tilde{\gamma}}{\Phi^2 - 1} \quad \text{for} \quad \Phi = \left(\frac{v_{\min} + \rho v_0 + \gamma\,\delta q}{(1+\rho)v_0}\right)^{\frac{\tilde{\gamma}}{\gamma}}, \tag{3.5}$$

where

$$v_{\min}^2 = \gamma^2\delta q^2 + 2\gamma\rho\delta q v_0 + v_0^2, \quad \delta q = \frac{k^{1-\beta} - F_0^{1-\beta}}{1-\beta} \quad \text{and} \quad \delta\tilde{q} = \frac{k^{1-\tilde{\beta}} - F_0^{1-\tilde{\beta}}}{1-\tilde{\beta}}. \tag{3.6}$$

The effective strike k is a *floored* initial strike: all the effective parameters formulae based on the heat-kernel expansion work only for positive strikes. In our experiments, we used $k = \max(K, 0.1\,F_0)$. Note that the initial value of the rate F_0 is considered to be positive. See the remark in the introduction for negative F_0.

The first-order correction is more complicated and is given by

$$\frac{\tilde{v}_0^{(1)}}{\tilde{v}_0^{(0)}} = \tilde{\gamma}^2\,\sqrt{1+\tilde{R}^2}\,\frac{\frac{1}{2}\ln\left(\frac{v_0 v_{\min}}{\tilde{v}_0^{(0)}\tilde{v}_{\min}}\right) - \mathcal{B}_{\min}}{\tilde{R}\,\ln\left(\sqrt{1+\tilde{R}^2} + \tilde{R}\right)} \quad \text{for} \quad \tilde{R} = \frac{\delta q\,\tilde{\gamma}}{\tilde{v}_0^{(0)}},$$

where $\tilde{v}_{\min} = \sqrt{\tilde{\gamma}^2\delta q^2 + \left(\tilde{v}_0^{(0)}\right)^2}$ and $\mathcal{B}_{\min}$ is the so-called parallel transport, defined as

$$\mathcal{B}_{\min} = -\tfrac{1}{2}\frac{\beta}{1-\beta}\frac{\rho}{\sqrt{1-\rho^2}}\left(\pi - \arccos\left(-\frac{\delta q\,\gamma + v_0\,\rho}{v_{\min}}\right) - \arccos\rho - I\right)$$

and

$$I = \begin{cases} \frac{2}{\sqrt{1-L^2}} \left(\arctan \frac{u_0+L}{\sqrt{1-L^2}} - \arctan \frac{L}{\sqrt{1-L^2}} \right) & \text{for } L < 1, \\ \\ \frac{1}{\sqrt{L^2-1}} \ln \frac{u_0\left(L+\sqrt{L^2-1}\right)+1}{u_0\left(L-\sqrt{L^2-1}\right)+1} & \text{for } L > 1, \end{cases} \tag{3.7}$$

where

$$L = \frac{v_{\min}(1-\beta)}{k^{1-\beta}\,\gamma\,\sqrt{1-\rho^2}} \quad \text{and} \quad u_0 = \frac{\delta q\,\gamma\rho + v_0 - v_{\min}}{\delta q\,\gamma\sqrt{1-\rho^2}}.$$

See also [9] and [13].

At the end of this section, we comment on the arbitrage-free property of the free-boundary SABR model. As a *real* process, the free-boundary SABR model is naturally arbitrage-free. On the other hand, its approximation described in this subsection is, strictly speaking, not arbitrage-free (except in the case of the zero correlation when the approximation becomes exact). However, in our efficient volatility expansion $\tilde{v}_0$ we use the *fixed* efficient strike $k = 0.1\,F_0$ for $K < 0.1\,F_0$. This means that prices for such strikes are given by the zero-correlation model with *strike-independent* coefficients and, consequently, are arbitrage-free. For the other strikes, given high approximation accuracy, we can call the resulting analytical formula *quasi*-arbitrage-free.

4 Numerical Experiments

Calibration to real data. We start with a real data example of a 1Y15Y CHF swaption from 10-Feb-2015 with a forward rate of $F_0 = 0.56\%$. The swaption prices are quoted in terms of normal implied volatility (bps). We calibrate free-boundary and shifted SABR models to this data by using our analytical approximations. The output is presented in Table 1 and graphed in Figure 4.

Strike	Target	free-bdry	shifted
0.06%	23.5	23.5	24.6
0.31%	44.7	44.5	43.3
0.56%	59.3	59.2	58.7
0.81%	71.7	71.7	71.8
1.06%	83.0	83.1	82.9
1.56%	103.5	103.8	103.6
2.56%	140.4	140.2	140.7

Table 1: Normal implied volatilities (in bps) for the calibrated free-boundary SABR and shifted SABR models.

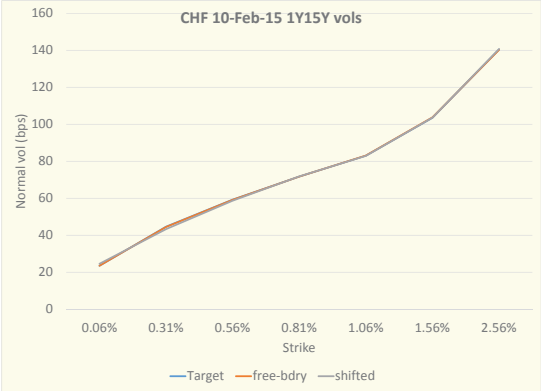

Figure 4: Plot of the free-boundary normal implied volatilities (bps) from Table 1.

The calibration errors for the free-boundary SABR model and the shifted SABR model are shown in Figure 5. We see that the calibration error for the free-boundary SABR is small. Table 2 shows the calibrated values of $\alpha = v_0$, ρ, γ, and β. The value of the shift for the shifted SABR model is 2%.

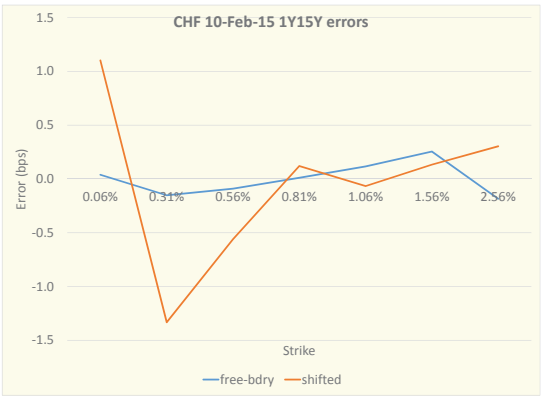

Figure 5: Calibration errors (in bps) for free-boundary and shifted SABR models. Wee see that the free-boundary SABR model has less calibration error.

Param	free SABR	shifted SABR
α	0.051	0.011
β	0.417	0.167
ρ	0.990	0.999
γ	0.658	1.080

Table 2: Calibrated parameters for the free-boundary SABR and shifted SABR models.

Note the extremely high values of the correlation ρ and fairly high values of the vol-of-vol γ. The reason for such high correlation is a very steep skew prevailing in the current CHF.

We now study the *accuracy* of the analytical approximation for the free SABR model. First, let us briefly address the Monte Carlo simulation scheme. (See [5] for more details.) Suppose that we have simulated the stochastic volatility for all timesteps and paths v_t.[3] Our goal is to simulate $F_{t+\Delta t}$ given

[3]This is trivial for the lognormal process.

this information. The first thing to try is an Euler scheme without any boundary condition (i.e., a free boundary), which is

$$F_{t+\Delta t} = F_t + |F_t|^\beta \, v_t \, \Delta W_1(t).$$

One can explicitly check that the Euler scheme does not work for points close to zero even when $\Delta t \to 0$. Instead, we need to come up with a more careful scheme based on a numerical inversion of the CDF, which can be found in [5]. However, such a procedure is very slow, and we prefer to come up with a regime-switching scheme similar to [1] in order to accelerate the simulations. For values away from the boundary, we use moment matching to approximate $F_{t+\Delta t}$ via the quadratic Gaussian step, while for near-boundary values, we numerically invert the CDF.

In Table 3, we compare the Monte Carlo simulations described above (Exact) and our analytical formula based on the map to the zero correlation SABR model (Analyt) for the calibrated parameters in Table 2. As shown, the approximation is excellent.

K	Analyt	Exact	Diff
0.06%	24	24	-0.7
0.31%	44	45	-0.7
0.56%	59	60	-0.8
0.81%	72	73	-0.8
1.06%	83	84	-0.9
1.56%	104	105	-1.0
2.56%	140	142	-1.3

Table 3: Comparison of implied volatilities from Monte Carlo simulations (Exact) and the analytical formula (Analyt) with the parameters from Table 2.

Approximation accuracy analysis. We analyze the approximation accuracy for two additional sets of inputs. These input data sets are negatively correlated, which is a more classical situation. The input data that is used is shown in Table 4, and the implied volatilities from the Monte Carlo simulation method are compared against the analytical approximation for both sets of inputs in Table 5.

Parameter	Symbol	Value for Input I	Value for Input II
Rate Initial Value	F_0	50 bps	1%
SV Initial Value	v_0	$0.6\, F_0^{1-\beta}$	$0.3\, F_0^{1-\beta}$
Vol-of-Vol	γ	0.3	0.3
Correlations	ρ	-0.3	-0.3
Skews	β	0.25	0.25
Maturities	T	3Y	10Y

Table 4: Setups for the free-boundary SABR model.

	Input I			Input II		
K	**Analyt**	**Exact**	**Diff**	**Analyt**	**Exact**	**Diff**
-0.95	30.87	30.93	-0.06	40.05	40.86	-0.81
-0.8	29.83	29.95	-0.12	38.43	39.24	-0.81
-0.65	28.80	28.97	-0.17	36.80	37.60	-0.80
-0.5	27.79	27.99	-0.20	35.18	35.97	-0.78
-0.35	26.83	27.04	-0.21	33.59	34.33	-0.74
-0.2	25.95	26.15	-0.20	32.05	32.73	-0.68
-0.05	25.30	25.46	-0.16	30.67	31.25	-0.58
0.1	25.77	25.85	-0.08	30.20	30.63	-0.43
0.25	26.63	26.69	-0.06	30.19	30.51	-0.31
0.4	27.33	27.39	-0.06	30.14	30.41	-0.27
0.55	27.90	27.97	-0.06	30.06	30.31	-0.25
0.7	28.38	28.45	-0.07	30.00	30.22	-0.23
0.85	28.80	28.87	-0.07	29.98	30.18	-0.20
1	**29.18**	**29.25**	**-0.07**	**30.05**	**30.22**	**-0.17**
1.15	29.53	29.60	-0.07	30.24	30.36	-0.12
1.3	29.87	29.94	-0.07	30.56	30.63	-0.07
1.45	30.22	30.29	-0.06	31.03	31.04	-0.01
1.6	30.58	30.63	-0.06	31.63	31.58	0.04
1.75	30.95	30.99	-0.05	32.35	32.26	0.09
1.9	31.33	31.37	-0.04	33.17	33.04	0.13

Table 5: Differences in implied volatilities (in bps) between simulations (Exact) and analytics (Analyt). The bold line ($K = 1$) represents the ATM strike.

The implied volatility results that are shown in Table 5 are plotted in Figure 6

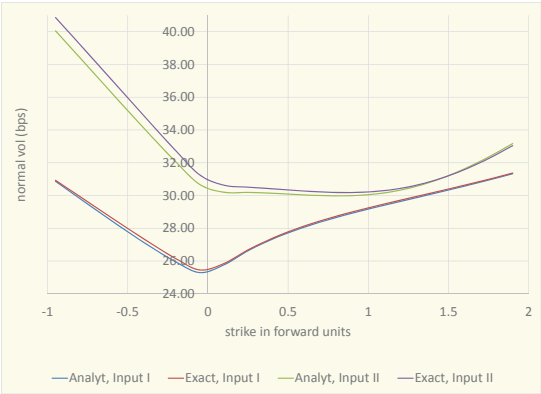

Figure 6: Plot of implied volatility for Monte Carlo simulation (Exact) and our method (Analyt).

The approximation is excellent for positive strikes for the 3Y data and for $K > \frac{1}{2}F_0$ for the 10Y data. For other strikes, we observe slight degeneration. We see that the normal implied volatility exhibits significant smiles, with the bottom between zero and the ATM strike. In general, increasing the volatility-of-volatility and maturity moves the vertex of the smile to the ATM strike.

5 Conclusion

We presented a natural generalization of the SABR model to negative rates, which is very important in the current low-rate environment, and described its properties. We derived an exact formula for the option price in the zero-correlation case and an efficient approximation for the general correlation case written in terms of a one-dimensional integral of elementary functions. The simplicity of the approximation allows for straightforward implementation. Moreover, the main formulae from our "absorbing" (standard) SABR approximation can be directly reused. Finally, we have numerically checked the accuracy of the approximation for option pricing.

Acknowledgments

The authors are indebted to Serguei Mechkov for discussions and help with the numerical implementation as well as to their other colleagues at Numerix, especially to Gregory Whitten and Serguei Issakov for supporting this work and to Nic Trainor for the excellent editing.

References

[1] Andersen, L. B. Efficient simulation of the heston stochastic volatility model. Available at http://ssrn.com/abstract=946405, 2007. Accessed 2014-10-03.

[2] Andreasen, J., and Huge, B. Expanded forward volatility. *Risk Magazine* (January 2013).

[3] Antonov, A., Konikov, M., Rufino, D., and Spector, M. Exact solution to CEV model with uncorrelated stochastic volatility, 2014. Available at SSRN.

[4] Antonov, A., Konikov, M., and Spector, M. SABR spreads its wings. *Risk Magazine* (August 2013), 58–63.

[5] Antonov, A., Konikov, M., and Spector, M. The free boundary SABR: Natural extension to negative rates. Available at http://ssrn.com/abstract=2557046, 2015. Accessed 2015-02-27.

[6] Balland, P., and Tran, Q. SABR goes normal. *Risk Magazine* (May 2013).

[7] Hagan, P. S., Kumar, D., Lesniewski, A. S., and Woodward, D. E. Managing smile risk. *Wilmott Magazine 1* (November 2002), 84–108.

[8] Hagan, P. S., Kumar, D., Lesniewski, A. S., and Woodward, D. E. Arbitrage free SABR. *Wilmott Magazine* (January 2014).

[9] Henry-Labordere, P. *Analysis, Geometry, and Modeling in Finance: Advanced Methods in Option Pricing*. Chapman & Hall, 2008.

[10] Islah, O. Solving SABR in exact form and unifying it with LIBOR market model. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1489428009, 2009.

[11] Jeanblanc, M., Yor, M., and Chesney, M. *Mathematical Methods for Financial Markets*. Springer, 2009.

[12] Mercurio, F. Pricing inflation-indexed derivatives. *Quantitative Finance* (2005), 289–302.

[13] Paulot, L. Asymptotic implied volatility at the second order with application to the SABR model. Available at http://ssrn.com/abstract=1413649, 2009.

[14] Rebonato, R., McKay, K., and White, R. *The SABR/LIBOR Market Model: Pricing, Calibration and Hedging for Complex Interest-Rate Derivative*. John Wiley & Sons Ltd, 2009.

Martingale and Distribution Tests for the Libor Market Model

Alexander Antonov* Dan Li† Leonard Tchuindjo‡

Abstract

Using martingale normalization, we find appropriate adjustment factors to simulated Libor rates, which we use for two purposes. First, we show that the adjusted model-simulated forward Libor rate is a martingale under the simulation measure. This allows us to construct a martingale restriction test for the Libor market model, which requires that in an arbitrage-free market, the expectation of each adjusted forward Libor rate simulated by a model calibrated to option-embedded instruments equals the corresponding forward Libor rate implied from the yield curve. This has practical implications for arbitrage trading. Second, we use the adjustment factors to construct a distribution test for the simulated forward Libor rates. We show how, in the simulated measure, the adjustment factors transform the distribution of each simulated Libor rate to match the initial distribution assumption of its corresponding forward Libor rate in its relevant measure. This has practical implications for model testing.

1 Introduction

The Libor market model (LMM) framework was independently proposed by Brace, Gatarek and Musiela [2] and Sandmann and Sonderman [13] and further improved by Jamshidian [9]. This framework is attractive because it can directly model observable rates such as Libor or swap rates, which is a significant improvement over short-rate models (e.g., Hull and White [8]) and forward-rate models (e.g., Heath, Jarrow, and Morton [10]), which model unobservable instantaneous rates.

The setup of the standard LMM is straightforward. We assume that each forward Libor rate follows a stochastic process with a known distribution under its relevant measure. However, because all rates are simultaneously simulated under a single measure, the stochastic differential equation (SDE) of each rate must be reformulated in this measure, which adds a time-dependent drift.[1] As a result, each forward Libor rate follows a stochastic process with a distribution that can be unknown in the simulating measure (especially if the initial process is lognormal). Market practitioners therefore face the problem of recovering the model distribution for the simulated forward Libor rate. For example, if the model setup assumes that each forward Libor is lognormal in its relevant measure, then one has to find a way to adjust the distribution of each simulated forward Libor so that it becomes lognormally distributed in the simulation measure. Constructing this adjustment allows us to test whether the distribution of the simulated forward Libor rates is concordant with the initial distribution assumption of the Libor market model.

Finding appropriate adjustments for the simulated forward Libor rate can also help verify the martingale restriction for the Libor rate market. This means that in the absence of arbitrage, the forward Libor rate simulated by the LMM calibrated to option-embedded market instruments (such as caps, floors, and swaptions) equals the corresponding forward rate implied from the yield curve constructed from Libor-based instruments (such as cash Libor rates, forward rate agreements, Eurodollar futures,

*Numerix, SVP of Quantitative Research, antonov@numerix.com

†Numerix, SVP of Financial Engineering, dli@numerix.com

‡Numerix, Director of Financial Engineering, ltchuindjo@numerix.com

[1] If the simulation is conducted in a measure associated with a particular rate, then the SDE for that rate under the simulation measure will have zero drift.

and swaps). The martingale restriction test can be conducted by using any option-pricing model to test for the absence of arbitrage in the corresponding market. Early studies were done by Longstaff [11], who conducted a test to check whether the underlying S&P value implied from call options on the S&P 100 through the model equals its actual market value. In [3], Brenner and Eom conducted the test on call and put options on the S&P 500. Busch [4] has also tested the martingale restriction on the USD-GBP call and put currency market. These authors conducted their analysis in the Black-Scholes framework. However, no such test has been done for Libor rates to the best of our knowledge.

We use martingale normalization to find appropriate adjustment factors, which are built upon a proposed sample stochastic discount factor, and show that when the simulated forward Libor rate is adjusted by these factors, its expectation under the simulating measure equals the value of the corresponding forward Libor rate implied from the yield curve. We set up a one-factor lognormal LMM and we check this condition with USD market data for 31 December 2014. We find that the adjusted forward three-month Libor rates simulated by the LMM calibrated to swaptions are very close to the forward three-month Libor rates implied from the Eurodollar futures and swap data. This shows that there are very limited arbitrage opportunities and negligible frictions in the USD Libor market.[2] Furthermore, we use the above factors to adjust the distributions of the simulated forward Libor rates so that they become lognormal. We thus recover the initial lognormal assumption of our LMM setup. A Kolmogorov-Smirnov test supports this finding. However, it is important to note that our proposed martingale and distribution tests are not based on a specific assumption for the distribution of the underlying Libor rate.

This paper is structured as follows. Section 2 provides a brief overview of the LMM framework. In Section 3, we propose a stochastic discount factor, which is used to construct the adjustment factors for the simulated forward Libor rates in Section 4, where we show that in the absence of arbitrage, the expectation of each adjusted simulated forward Libor rate equals the corresponding forward Libor rate implied from the yield curve. Numerical results illustrate this result. In Section 5, we prove that, in the simulated measure, the adjustment factors transform the distribution of each simulated Libor rate to match the initial distribution assumed for its corresponding forward Libor rate in the LMM setup. We use a Kolmogorov-Smirnov test to compare the distributions. Section 6 summarizes the results and proposes areas of further investigation.

2 An Overview of the LMM

Consider a discrete set of times $T = \{T_i\}_{i=0}^{M}$, where $0 = T_0 < T_1 < \cdots < T_M < \infty$. At any time $t \geq 0$, define $P(t, T_i)$ to be the time-t price of a zero-coupon bond maturing at $T_i \geq t$. We assume this zero-coupon bond to be default-free; i.e., $P(t, T_i) > 0$ for all $t \in [0, T_i]$. The forward Libor rate between times T_i and T_{i+1}, as seen at time $t \leq T_i$, can be defined as

$$L_i(t) \equiv \frac{\left(P(t, T_i) - P(t, T_{i+1})\right)/\delta_i}{P(t, T_{i+1})}, \tag{2.1}$$

where $\delta_i = T_{i+1} - T_i$ is the day-count fraction between T_i and T_{i+1}. In (2.1), the numerator is a tradable asset, as it is a linear combination of two tradable assets, and the denominator is strictly positive, as it is the price of a default-free bond. Thus, following Geman, Karoui and Rochet [6], the stochastic process of the forward Libor rate $\{L_i(t) : 0 \leq t \leq T_i\}$ is a martingale under the measure $\mathbb{Q}^{i+1}$ associated with the numeraire $\{P(t, T_{i+1}) : 0 \leq t \leq T_{i+1}\}$.

Because of the martingale property, the stochastic process of the forward Libor rate is driftless under $\mathbb{Q}^{i+1}$. Furthermore, by the martingale representation theorem, this stochastic process can be represented as a function of a finite number of Brownian motions. Therefore, if we assume this forward Libor rate to be lognormally distributed, we can write

$$dL_i(t) = L_i(t) \sum_{k=1}^{K} \sigma_{i,k}(t) dW_{i+1,k}(t), \tag{2.2}$$

[2] All numerical computations (e.g., yield curve construction, model calibration, and rate simulations) were done using Numerix CrossAsset.

where each $\sigma_{i,k} : \mathbb{R}_+ \to \mathbb{R}$ is bounded and square integrable on $[0, T_M]$, and $W_{i+1,k}(t)$ represents the time-t value of a $\mathbb{Q}^{i+1}$-standard Brownian motion. However, the process of this forward Libor will not be a martingale under a different measure. As in [14], one can model all forward Libor rates using the rolling spot measure, which is defined with the numeraire

$$N(t) \equiv P(t, T_{i+1}) \prod_{k=0}^{i} P(T_k, T_{k+1})^{-1} \tag{2.3}$$

for $t \in [T_i, T_{i+1}]$. In this measure, the stochastic differential equation of the forward Libor rate is

$$dL_i(t) = L_i(t) \sum_{j=\eta(t)}^{i} \left(\frac{\delta_j L_j(T_j)}{1 + \delta_j L_j(T_j)} \sum_{k=1}^{K} \sigma_{i,k}(t)\sigma_{j,k}(t) \right) dt + L_i(t) \sum_{k=1}^{K} \sigma_{i,k}(t), dW_{i+1,k}(t), \tag{2.4}$$

where $\eta(t) = m + 1$ for $t \in (T_m, T_{m+1}]$.

3 The Sample Stochastic Discount Factor

We consider a one-factor LMM with constant volatility, and we assume each forward Libor rate has a three-month tenor; i.e., δ_i is approximately three months.[3] In this case, the stochastic differential equation of the forward Libor rate in its relevant measure is

$$dL_i(t) = \sigma_i L_i(t) dW_{i+1}(t). \tag{3.1}$$

Let $T_0 = 0$ denote the current time. The current yield curve can be defined through the discount factor $D(0, t)$ for $t \in [0, T_M]$. In an N-path Monte Carlo simulation where the true measure is replaced by the sample measure S^*, we can construct a sample stochastic discount factor $A(0, t)$ such that its expectation equals the discount factor implied from the yield curve; i.e.,

$$\mathbb{E}^* [A(0, t)] = D(0, t), \tag{3.2}$$

where $\mathbb{E}^*$ represents the expectation operator under the sample measure S^*.[4]

Following a martingale normalization method as in [7], we describe how such a sample stochastic discount factor can be constructed. Let

$$B(0, T_i) = \prod_{j=0}^{i-1} (1 + \delta_j L_j(T_j))^{-1},$$

where the values $L_j(T_j)$ are simulated by using Equation (3.1). For maturity $T_i \geq 0$, the sample expectation $\mathbb{E}^* [B(0, T_i)]$ is computed as the arithmetic average of the results from all paths, $\frac{1}{N} \sum_{\text{paths}} B(0, T_i)$. Certainly, this sample expectation does not equal the discount factor for the same maturity; i.e., $B(0, t)$ does not satisfy (3.2). However, we can rescale $B(0, T_i)$ to force this condition to hold. Indeed, we wish to find a deterministic c_i such that

$$\mathbb{E}^* [c_i B(0, T_i)] = D(0, T_i).$$

As c_i is deterministic, we simply have

$$c_i = \frac{D(0, T_i)}{\mathbb{E}^* [B(0, T_i)]}. \tag{3.3}$$

Thus, Equation (3.2) is obtained if we set

$$A(0, T_i) = c_i B(0, T_i). \tag{3.4}$$

Table 1 in Appendix A shows a numerical illustration of Equation (3.2) for sample stochastic discount factors with maturities ranging from three months to five years with quarterly intervals. The sample stochastic discount factors were created from a 1000-path Monte Carlo simulation.

[3] This corresponds to setting "Volatility Type" to Flat Volatility and leaving "Correlation Type" empty in Numerix CrossAsset. (See, e.g., [12]).

[4] Numerix CrossAsset has an implementation of such a sample discount factor, called ADP. See, e.g., [5] for details.

4 The Martingale Test

For the forward Libor rate $L_i(t)$, Equation (3.1) holds only in the measure $\mathbb{Q}^{i+1}$. In fact, a drift term must be added to express (3.1) in any other measure. Therefore, the terminal value $L_i(T_i)$ of the simulation will not be a martingale under the sample measure. Moreover, its distribution might be unknown.

On the one hand, under the measure $\mathbb{Q}^{i+1}$, the expected value of the forward Libor rate at T_i equals the yield-curve-implied forward Libor rate maturing at T_i; i.e.,

$$\mathbb{E}^{\mathbb{Q}^{i+1}}\left[L_i(T_i)\right] = L_i(0). \tag{4.1}$$

One the other hand, define

$$\tilde{L}_i(T_i) = P(T_i, T_{i+1})L_i(T_i),$$

where $P(T_i, T_{i+1})$ is the simulated price at T_i of a zero-coupon bond maturing at T_{i+1}. Let

$$N^*(t) = \frac{1}{A(0,t)}$$

for any time $t \geq 0$. By the change of measure technique, we have

$$\mathbb{E}^*\left[\frac{\tilde{L}_i(T_i)}{N^*(T_i)}\right] N^*(0) = \mathbb{E}^{\mathbb{Q}^{i+1}}\left[\frac{\tilde{L}_i(T_i)}{P(T_i, T_{i+1})}\right] P(0, T_{i+1}). \tag{4.2}$$

As $N^*(0) = 1$ and $P(0, T_{i+1}) = D(0, T_{i+1})$, and using the definition of $\tilde{L}_i$, (4.2) can be rewritten as

$$\mathbb{E}^*\left[\frac{P(T_i, T_{i+1})L_i(T_i)}{N^*(T_i)D(0, T_{i+1})}\right] = \mathbb{E}^{\mathbb{Q}^{i+1}}\left[L_i(T_i)\right]. \tag{4.3}$$

Combining Equations (4.1) and (4.3) gives

$$\mathbb{E}^*\left[\omega_i L_i(T_i)\right] = L_i(0), \tag{4.4}$$

where

$$\omega_i \equiv A(0, T_i)\frac{P(T_i, T_{i+1})}{D(0, T_{i+1})}.$$

As $\omega_0 = 1$, Equation (4.4) can be rewritten as

$$\mathbb{E}^*\left[\omega_i L_i(T_i)\right] = \omega_0 L_i(0). \tag{4.5}$$

Thus, $\omega_i L_i(T_i)$ satisfies the martingale property under the measure S^*. Therefore, the adjusted simulated forward Libor rates are martingales under the sample measure. Table 2 in Appendix A shows a numerical illustration of this property for forward three-month Libor rates with maturities ranging from three months to five years. This table also shows the expectation of the simulated weights. The results come from a 1000-path Monte Carlo simulation.

5 The Distribution Test

Equation (4.5) shows how the simulated forward Libor rate $L_i(T_i)$ can be adjusted to become a martingale under the sample measure S^*. However, we still do not know its distribution in the sample measure. Recall the cumulative distribution function of $L_i(T_i)$ can be represented as

$$F_i(x) = \mathbb{E}^{\mathbb{Q}^{i+1}}\left[\mathbb{I}_{\{L_i(T_i)<x\}}\right], \tag{5.1}$$

where $x > 0$ and $\mathbb{I}_A$ represents the indicator function of the set A. Define

$$U_i(x) = P(T_i, T_{i+1})\mathbb{I}_{\{L_i(T_i)<x\}}. \tag{5.2}$$

By substituting $U_i(x)$ for $\tilde{L}_i(T_{i+1})$ in (4.2), we obtain

$$\mathbb{E}^*\left[\frac{U_i(x)}{N^*(T_i)}\right] N^*(0) = \mathbb{E}^{\mathbb{Q}^{i+1}}\left[\frac{U_i(x)}{P(T_i,T_{i+1})}\right] P(0,T_{i+1}). \tag{5.3}$$

Again, as $N^*(0) = 1$ and $P(0,T_{i+1}) = D(0,T_{i+1})$, we have

$$\mathbb{E}^*\left[\frac{U_i(x)}{N^*(T_i)D(0,T_{i+1})}\right] = \mathbb{E}^{\mathbb{Q}^{i+1}}\left[\frac{U_i(x)}{P(T_i,T_{i+1})}\right]. \tag{5.4}$$

Using the definitions of $U_i(x)$ and ω_i, as in (5.2) and Section 4 respectively, (5.4) becomes

$$\mathbb{E}^*\left[\omega_i \mathbb{I}_{\{L_i(T_i)<x\}}\right] = \mathbb{E}^{\mathbb{Q}^{i+1}}\left[\mathbb{I}_{\{L_i(T_i)<x\}}\right]. \tag{5.5}$$

Then Equations (5.1) and (5.5) give

$$\mathbb{E}^*\left[\omega_i \mathbb{I}_{\{L_i(T_i)<x\}}\right] = F_i(x). \tag{5.6}$$

Equation (5.6) shows that the factor ω_i adjusts the distribution of the simulated forward Libor $L_i(T_i)$ so that it becomes lognormally distributed in the simulation measure. To illustrate this result, we perform a Kolmogorov-Smirnov test on the adjusted distribution of each forward Libor rate simulated in Section 4. Note that for a 1000-point sample, as in our case, the Kolmogorov-Smirnov test has critical values of 0.051545, 0.043007, and 0.038580 for α levels of 1%, 5%, and 10%, respectively. Appendix A shows the test statistics for the Kolmogorov-Smirnov test described above. With an α of 10% for each maturity, we fail to reject the null hypothesis that the adjusted simulated Libor rate is lognormally distributed in the simulation measure.

Note that (5.6) does not depend on any particular form of the cumulative distribution function, i.e., on any assumed distribution. The result holds for any distribution of the underlying Libor rate.

6 Conclusion

In this paper, we construct adjustment factors that transform all LMM-simulated forward Libor rates into martingales in the simulation measure. We use a one-factor lognormal LMM and find that in the USD Libor market, the forward three-month Libor rates (up to five years) implied from option-embedded instruments equal the corresponding forward Libor rates implied from option-free instruments. This shows that there are almost no arbitrage opportunities in the USD Libor market. Small differences can be explained by frictions stemming from bid-ask spreads. These adjustment factors also allow us to transform the distributions of the simulated forward Libor rates to recover the initial lognormal distribution assumption of our LMM setup. This study can be extended in various directions, including using different volatility types, changing the initial distribution assumption of the forward Libor rates, or adding more factors to the LMM. Moreover, the test can be performed on other currency markets.

Acknowledgments

The authors thank Nic Trainor and Anna Basanskaya for their review and suggestions. The authors remain responsible for all errors and omissions.

Appendix A Numerical Results

Here we display the results of the tests we performed. Table 1 shows the results of testing condition (3.2), Table 2 shows the results of testing the martingale condition (4.5), and Table 3 shows the results of the Komogorov-Smirnov test for the lognormality of the adjusted simulated forward Libor rate. All tests were performed with 1000 Monte Carlo paths.

Date	T_i	$D(0,T_i)$	$\mathbb{E}^*\left[A(0,T_i)\right]$	$D(0,T_i) - \mathbb{E}^*\left[A(0,T_i)\right]$
31-Dec-14	0.00	1.0000000000	1.0000000000	0.0000000000
31-Mar-15	0.25	0.9993614081	0.9993614081	0.0000000000
30-Jun-15	0.50	0.9986095184	0.9986095184	0.0000000000
30-Sep-15	0.75	0.9973844368	0.9973844368	0.0000000000
31-Dec-15	1.00	0.9956225621	0.9956225621	0.0000000000
31-Mar-16	1.25	0.9932188573	0.9932188573	0.0000000000
30-Jun-16	1.50	0.9901739071	0.9901739071	0.0000000000
30-Sep-16	1.75	0.9865333822	0.9865333822	0.0000000000
30-Dec-16	2.00	0.9823806984	0.9823806984	0.0000000000
31-Mar-17	2.25	0.9777345261	0.9777345261	0.0000000000
30-Jun-17	2.50	0.9727683334	0.9727683334	0.0000000000
29-Sep-17	2.75	0.9674839017	0.9674839017	0.0000000000
29-Dec-17	3.00	0.9619691093	0.9619691093	0.0000000000
30-Mar-18	3.25	0.9562668430	0.9562668430	0.0000000000
29-Jun-18	3.50	0.9505420559	0.9505420559	0.0000000000
28-Sep-18	3.75	0.9448515408	0.9448515408	0.0000000000
31-Dec-18	4.00	0.9390091940	0.9390091940	0.0000000000
29-Mar-19	4.24	0.9331238323	0.9331238323	0.0000000000
28-Jun-19	4.49	0.9270659080	0.9270659080	0.0000000000
30-Sep-19	4.75	0.9208495638	0.9208495638	0.0000000000
31-Dec-19	5.00	0.9148058463	0.9148058463	0.0000000000

Table 1: Sample stochastic discount factor test results.

Date	T_i	$\mathbb{E}^*\left[\omega_i\right]$	$L_i(0)$	$\mathbb{E}^*\left[\omega_i L_i(T_i)\right]$	$L_i(0) - \mathbb{E}^*\left[\omega_i L_i(T_i)\right]$
31-Dec-14	0.00	1.0000	0.0026	0.0026	0.0000
31-Mar-15	0.25	1.0000	0.0030	0.0030	0.0000
30-Jun-15	0.50	1.0000	0.0048	0.0048	0.0000
30-Sep-15	0.75	1.0000	0.0069	0.0069	0.0000
31-Dec-15	1.00	1.0000	0.0096	0.0096	0.0000
31-Mar-16	1.25	1.0000	0.0122	0.0122	0.0000
30-Jun-16	1.50	1.0000	0.0144	0.0144	0.0000
30-Sep-16	1.75	1.0000	0.0167	0.0167	0.0000
30-Dec-16	2.00	1.0001	0.0188	0.0188	0.0000
31-Mar-17	2.25	1.0000	0.0202	0.0202	0.0000
30-Jun-17	2.50	1.0000	0.0216	0.0216	0.0000
29-Sep-17	2.75	1.0000	0.0227	0.0227	0.0000
29-Dec-17	3.00	1.0001	0.0236	0.0236	0.0000
30-Mar-18	3.25	1.0000	0.0238	0.0238	0.0000
29-Jun-18	3.50	1.0000	0.0238	0.0238	0.0000
28-Sep-18	3.75	1.0002	0.0238	0.0238	0.0000
31-Dec-18	4.00	1.0000	0.0258	0.0258	0.0000
29-Mar-19	4.24	1.0000	0.0259	0.0259	0.0000
28-Jun-19	4.49	1.0000	0.0259	0.0259	0.0000
30-Sep-19	4.75	1.0001	0.0259	0.0259	0.0000
31-Dec-19	5.00	1.0000	0.0269	0.0269	0.0000

Table 2: LMM martingale test results.

Date	T_i	Test Statistics
31-Mar-15	0.25	0.002316
30-Jun-15	0.50	0.004037
30-Sep-15	0.75	0.004519
31-Dec-15	1.00	0.004379
31-Mar-16	1.25	0.010219
30-Jun-16	1.50	0.009851
30-Sep-16	1.75	0.019725
30-Dec-16	2.00	0.009038
31-Mar-17	2.25	0.006100
30-Jun-17	2.50	0.010026
29-Sep-17	2.75	0.014423
29-Dec-17	3.00	0.007934
30-Mar-18	3.25	0.009531
29-Jun-18	3.50	0.011102
28-Sep-18	3.75	0.013060
31-Dec-18	4.00	0.012425
29-Mar-19	4.24	0.015087
28-Jun-19	4.49	0.016621
30-Sep-19	4.75	0.018845
31-Dec-19	5.00	0.019678

Table 3: Kolmogorov-Smirnov test statistic.

References

[1] Black, F., and Scholes, M. The Pricing of Options and Corporate Liabilities. *The Journal of Political Economy 81*, 3 (May–June 1973), 637–654.

[2] Brace, A., Gatarek, D., and Musiela, M. The market model of interest rate dynamics. *Mathematical Finance 7*, 2 (1996), 127–154.

[3] Brenner, M., and Eom, Y. No-arbitrage option pricing: New evidence on the validity of the martingale property, 1997. Working Paper Series.

[4] Busch, T. Testing the martingale restriction for option implied densities. *Review of Derivatives Research 11* (2008), 61–81.

[5] Eliezer, D., Issakov, S., Mechkov, S., and Trainor, N. Numerix Monte Carlo convergence. Numerix technology papers, Numerix, 2014.

[6] Geman, H., Karoui, N. E., and Rochet, J. Changes of numeraire, changes of probability measure and option pricing. *Journal of Applied Probability 32*, 2 (1995), 443–458.

[7] Glasserman, P. *Monte Carlo Methods in Financial Engineering*. Springer, 2003.

[8] Hull, J., and White, A. Pricing interest rate derivative securities. *The Review of Financial Studies 3* (1990), 573–592.

[9] Jamshidian, F. Libor and swap market models and measures. *Finance and Stochastics 1* (1997), 293–330.

[10] Jarrow, R., Heath, D., and Morton, A. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica 60* (1992), 77–105.

[11] Longstaff, F. A. Option pricing and the martingale restriction. *The Review of Financial Studies 8*, 4 (1995), 1091–1124.

[12] Mihai, I. LMM specifications in Numerix CrossAsset. Numerix support papers, Numerix, 2013.

[13] Miltersen, K., Sandmann, and Sondermann, D. Closed-form solutions for term structure derivatives with log-normal interest rates. *Journal of Finance 52*, 1 (1997), 409–430.

[14] Numerix. Models in Numerix analytics, 2015.

"Hot-Start" Initialization of the Heston Model

Serguei Mechkov*

Abstract

We suggest a new way of setting up multifactor models with hidden variables. We claim that the standard initial condition, which assigns some fixed value to the stochastic volatility subprocess, is illogical and greatly underestimates the effect of the hidden variable. For instance, a stochastic volatility model generates a significantly weaker implied volatility smile at short maturities. A good initial condition should specify the distribution of the hidden variable instead of a particular fixed value. The most straightforward way of initializing a hidden variable is by specifying its equilibrium distribution, which assumes that this component of the multifactor process has been started well before the observable part. As a practical example, the Heston model is considered.

1 Introduction

A logical way of making a model more closely fit observed behavior is to augment its dimension by adding a hidden process that somehow affects the evolution of the principal observable value. The complete model specification then acquires the parameters of the hidden process and its possible correlation with the principal process. This also requires describing how the hidden process starts, which is the main subject of this paper.

The very fact that the process is hidden makes it illogical to assume that its latent variable has some definite value today as part of the initial conditions of the model. A more logical initial condition is a distribution of the latent variable based on the previous history of the market. We will refer to this as a "hot-start" initialization of the process. The market history for the principal observable value may be known and somehow taken into account or it may be ignored so that the distribution of the latent variable is simply an equilibrium achieved by the process started some time ago. The parameters of the implied historical process may be a continuation of the parameters used for the future evolution or independently adjusted by some calibration. The main point is that the initial condition for the latent variable should not just be some fixed value, but instead must be a hot-start distribution.

These considerations have already been taken into account by Dragulescu and Yakovenko [2] in the econometric context for quantitative comparisons of modeled stock distributions against historical records. Although not totally unnoticed, the hot-start approach seems to have been essentially ignored by successors (and sometimes even explicitly disregarded as irrelevant [1]). In particular, we are not aware of any attempt to apply it to pricing market instruments. In this domain, the latent process is always initialized by a single value that is somehow selected. Even if this value is calibrated by a best fit to market data, it is still just a number. As a result, the stochasticity of the latent process is suppressed at short maturities and only takes effect after a certain relaxation time. In our opinion, hot-start initialization provides a method to better fit the market at all maturities.

2 Reparameterized Heston Model

Below we concentrate on the stochastic volatility latent process. A typical example is the Heston process, which is the most convenient and popular model of this sort. The original formulation [5] considers a

*Numerix, SVP Quantitative Research, meshkov@numerix.com

geometric Brownian motion for the stock S and then postulates that the squared instantaneous volatility of $x = \log S$ (denoted v for "variance") follows the mean-reverting Cox-Ingersoll-Ross (CIR) process:

$$\begin{aligned} dx_t &= r dt - \frac{1}{2} v_t dt + \sqrt{v_t} dW_x, \\ dv_t &= a\left(\theta - v_t\right) dt + \varkappa \sqrt{v_t} dW_v, \\ \langle dW_x dW_v \rangle &= \rho dt. \end{aligned} \tag{2.1}$$

Leaving aside the trivial deterministic drift r, which is due to market discounting, the Heston process has four parameters: the long-term mean $\theta > 0$, the reversion speed $a > 0$ of the variance, the volatility $\varkappa > 0$ of the variance, and the correlation $-1 \le \rho \le 1$ between the two Brownian motions W_x and W_v.

The presence of the stochastic volatility allows the Heston model to adequately capture the implied volatility smile of the actual market at longer maturities. It is known, however, (see, e.g., Gatheral [4]) that this smile is unrealistically weak at short maturities. This apparently fundamental property usually prompts the conclusion that the stochastic volatility approach is intrinsically incapable of producing strong volatility smiles at short maturities, and therefore the model must include jumps. Such an extension (e.g., the Bates model) formally fixes the problem, but it also introduces new complexity into the calibration because of the need to determine the additional parameters related to the jump process.

We stress that the weak smile of the Heston process is not an artifact of the process itself, but rather reflects the essentially incorrect initialization of the stochastic volatility component by a single initial value v_0. By analyzing the effect of the appropriate hot-start initialization, we show that it can generate very strong implied volatility smiles at short maturities, suggesting that the Heston model is actually more powerful in fitting the market than is usually assumed.

The classic Heston parameterization is inconvenient in several respects. The most important problem is its very poor time-dependent behavior, which only allows very gradual time evolution of the implied volatility. For flexibility of extensions and a more intuitive connection of the parameters to the implied volatility surface, we abandon the classic formulation of the Heston model (2.1) and switch to the mathematically equivalent but more effective parameterization

$$\begin{aligned} dx_t &= r dt - \frac{1}{2} z_t \sigma^2 dt + \sigma \sqrt{z_t} dW_x, \\ dz_t &= a\left(\left(1 - z_t\right) dt + \gamma \sqrt{z_t} dW_z\right), \end{aligned} \tag{2.2}$$

where the CIR process drives a dimensionless stochastic multiplier z_t that reverts to 1. This multiplier enters the Black-Scholes-like diffusion for the log-return $x_t = \log(S_t)$ as a scaling factor for the explicit volatility parameter σ.

Here the speed of reversion a and the correlation ρ between the Brownian motions W_x and W_z are the same as in the classic formulation. The long-term mean θ is replaced by the volatility of the stock, with $\sigma = \sqrt{\theta}$, and instead of the volatility of variance $\varkappa$, we use the relative volatility of the stochastic factor, $\gamma = a^{-1}\theta^{-1/2}\varkappa$. In the following, we shall refer to γ as the *volatility noise*. The volatility noise actually controls the implied volatility smile at long maturities, whereas the correlation ρ affects the asymmetry (skew) of the profiles. The convenience of such scaling has been demonstrated in our previous work [7], where we show that the reparameterized Heston model remains valid even after taking the reversion a to infinity.

We assume that the evolution starts at $t_0 = 0$ from some arbitrary value x_0. The standard "cold-start" setup also includes the initial value for the stochastic variance factor z_0, the simplest default choice being $z_0 = 1$. An example of the implied volatility surface is presented in Figure 1. As expected, at short maturities the smile is weak. This is because the fluctuations of the volatility are not fully active during the initial relaxation period, which is of the order $a^{-1} = 1$ here.

3 Choice of the Hot-Start Distribution

The correct choice of the initial distribution of the latent process is vital to the hot-start concept. Dragulescu and Yakovenko [2] simply postulated zero knowledge of the current state of latent process

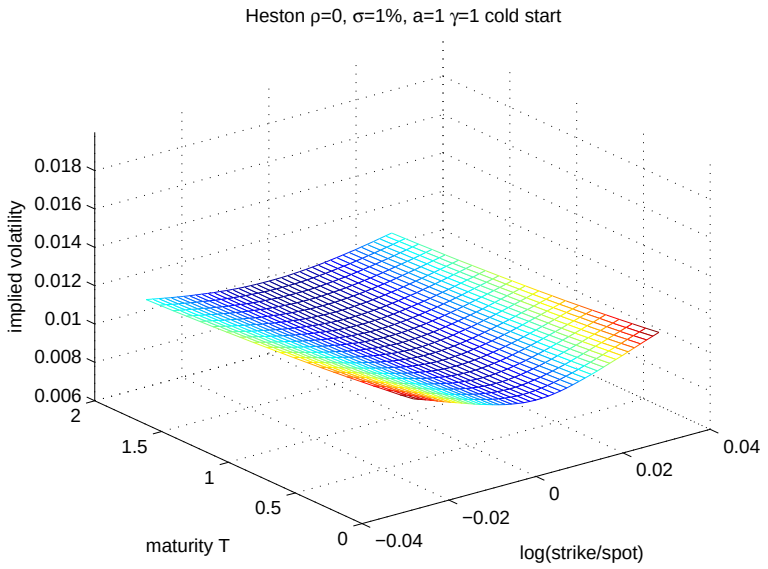

Figure 1: Implied volatility surface for a cold-start ($z_0 = 1$) Heston model.

for the evolution of the stochastic multiplier z, assuming that it has actually started much earlier and reached its equilibrium Gamma distribution with the density

$$F_{a,\gamma}(z) = \frac{\alpha^\alpha}{\Gamma(\alpha)} z^{\alpha-1} \exp\left[-\alpha z\right] , \tag{3.1}$$
$$\alpha = \frac{2}{a\gamma^2}.$$

by the time $t_0 = 0$. This straightforward choice may be quite a practical approximation in many cases. However, a more rigorous setup is also worth discussing.

In reality, some indirect information about the latent part of the process is available through the known history of the observable path because these components are coupled. Thus, we can expect that the hot-start distribution may be shifted and somewhat squeezed with respect to the equilibrium one. The question is how to quantify this effect. In the spirit of standard calibration, we are allowed to assume that today's market absorbs historical information and translates it into quotes. Under such an assumption, it is enough to somehow parameterize today's distribution of the latent factor and to calibrate these parameters together with those of the process.

Since the estimation cannot be very accurate in general, it should be sufficient to use only a couple of parameters related to the moments of the distribution. The "historical" effects we are trying to adapt to are not able to drastically change the principal features of the distribution. It is therefore most suitable to use the distribution from the same family as the equilibrium one.

For the Heston model, we suggest the Gamma distribution

$$F(z) = \frac{z^{\alpha-1}}{\Gamma(\alpha)\beta^\alpha} \exp\left[-\frac{z}{\beta}\right] \tag{3.2}$$

with two parameters: the shape α and the scale β. The moment generating function is

$$\mathbb{E}[e^{qz}] = (1 - \beta q)^{-\alpha} ,$$

and the equilibrium values

$$\alpha = \beta^{-1} = \alpha_0 = \frac{2}{a\gamma^2}$$

correspond to the unit average of z,

$$\mathbb{E}[z] = 1,$$

and the variance

$$\mathbb{E}[(z - \bar{z})^2] = \alpha\beta^2 = \alpha_0^{-1} = \frac{a\gamma^2}{2}.$$

We introduce two scaling parameters ζ and ω, determined by

$$\alpha = \frac{\alpha_0}{\omega},$$
$$\beta = \frac{\zeta\omega}{\alpha_0},$$

so that the average value of z and its variance become

$$\mathbb{E}[z] = \alpha\beta = \zeta$$

and

$$\mathbb{E}[(z - \bar{z})^2] = \alpha\beta^2 = \frac{\omega\zeta^2}{\alpha_0} = \omega\zeta^2\frac{a\gamma^2}{2},$$

respectively. The idea of this choice of scaling is that the bare equilibrium setting is obtained simply by setting $\zeta = \omega = 1$, and adjusting the average value of z through the ζ parameter does not perturb the shape α of the distribution.

When the model is applied in the econometric context, we do not usually rely on today's market quotes. Instead, we have to explicitly analyze the historical time series by considering the latent distribution to be a result of a long evolution, conditional on the actually recorded history of the observable variables. An accurate determination of the conditional distribution is a very challenging task. In fact, it is very close to the maximum likelihood estimation of multifactor process parameters from the historical market observables. It cannot be accomplished exactly, but there are workable approximations based on a recursive forward-in-time buildup of a few (e.g., two) of the lowest moments of the distribution. Once the recursion reaches the end (today), the hot-start distribution can be adjusted to match the obtained moments. In the case of the Heston process, the derivations are available, for instance, from Hurn et al. [6] and references therein. The parameterization introduced above is naturally valid for this purpose as well.

Computational Aspects

The complication introduced by the hot-start initialization to any computational framework is minimal because the underlying stochastic process remains the same.

The most straightforward is an extension of Monte Carlo simulation. Instead of initializing all the simulation paths by the same fixed value of the latent variable, the initial value should be sampled from the adjusted distribution. Essentially, this is equivalent to inserting a special Monte Carlo step at today's time node. The statistical accuracy of the pricing should not deteriorate.

In a finite difference scheme, the price is usually obtained by taking the payoff at the spot node of the lattice after back propagation to today's time node. For the hot-start evaluation, the payoff must be collected from the entire projection of the latent variable grid with the given spot value of the observable variable and averaged with the adjusted distribution.

Analytic pricing of European options is available for the Heston model. The price is usually obtained as a Fourier representation through the characteristic function of the marginal distribution of the stock at the option maturity date. As a function of the spot value z_0 of the latent stochastic multiplier, the characteristic function is an exponential of a linear form with coefficients determined by an analytically solvable Riccati equation. (For derivations, see, e.g., Dragulescu and Yakovenko [2] and, in our notations, [7].) The hot-start initialization just replaces the exponential by its integral with the Gamma distribution (3.2). While we do not write the integrand out here, it is clear that the numerical integration workload remains essentially the same.

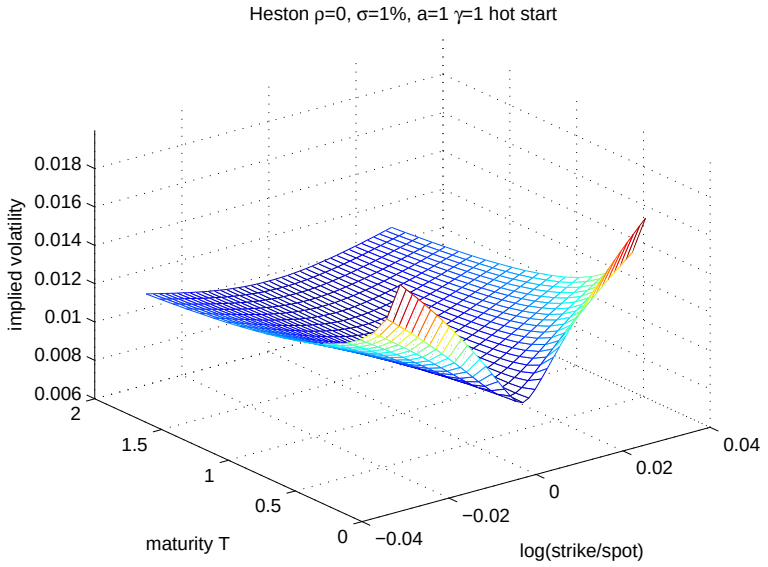

Figure 2: Implied volatility surface for a Heston model with a hot-start initialization from the equilibrium distribution (3.1).

As an example, for the same process parameters as those used for Figure 1, we applied the hot-start initialization with the full equilibrium distribution (3.1) and obtained the implied volatility surface shown in Figure 2. As expected, the surface changes only slightly on the horizons after the relaxation period a^{-1}, but the smile gets stronger at early maturities and demonstrates a tendency to diverge at very short ones. This is in striking contrast to the cold-start case, where implied volatility remains finite in the limit of zero time to maturity $T \to 0$. (See, e.g., [3].)

Note also that the short-maturity smile depends on the reversion speed a. When the reversion spreed increases, both the cold-start and the hot-start volatility surfaces display more extreme smiles at short maturities and become indistinguishable in the limit $a \to \infty$. This behavior was explored in [7]. Figure 3 provides a plot of the implied volatility of a such a fast-reversion Heston model.

Considering the short maturity asymptotic quantitatively, we notice that this limit reduces to pricing European options by the Black-Scholes model that corresponds to the first SDE in (2.2) with a fixed value of the factor z and averaging the result over the hot-start distribution (3.2). For a call option on the stock with spot value S_0 and strike $K = S_0 \exp[X + rT]$, the Black-Scholes price $C_z(X)$ is

$$S_0^{-1} C_z(X) = \frac{e^{-rT}}{2\pi} \frac{1}{\sqrt{2\pi V}} \int_X^{+\infty} \exp\left[-\frac{\left(x + \frac{1}{2}\sigma^2 z T\right)^2}{2\sigma^2 z T} \right] \left(e^x - e^X\right) dx.$$

In the limit as $T \to 0$, it is enough to work in the main exponential approximation for the Black-Scholes out-of-the-money price

$$S_0^{-1} C_z(X) \approx \exp\left[-\frac{X^2}{2\sigma^2 z T} \right]$$

and for the tail of Gamma distribution

$$F(z) \approx e^{-z/\beta}.$$

The resulting average price of the call option

$$C(X) \approx S_0 \exp\left[-\frac{X}{\sigma\sqrt{2T\beta}} \right]$$

is to be compared against the price of the same option under the Black-Scholes model with the implied

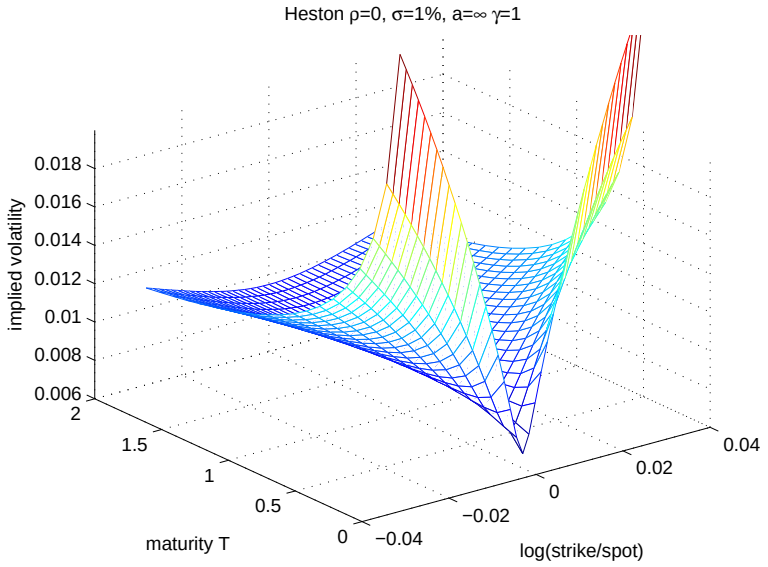

Figure 3: Implied volatility surface for a fast-reversion Heston model ($a = \infty$).

volatility $\tilde{\sigma}$,

$$C(X) \approx S_0 \exp\left[-\frac{X^2}{2\tilde{\sigma}^2 T}\right].$$

A match is obtained for

$$\tilde{\sigma}^2 = \sigma\sqrt{\frac{\beta}{2T}}\,|X|\,,$$

where we included the symmetric result for in-the-money strikes.

We see that the implied volatility indeed diverges at short maturities as $T^{-1/4}$. In addition, we see that, like the fast-reversion limit of the Heston model, the hot-start volatility surface is symmetric on short maturities. This is natural, because the hot-start distribution effectively comes from the past and thus is not correlated with subsequent stochastic moves of the stock. Note that only the leading term is symmetric and the overall smile is significantly skewed even at rather short maturities, as demonstrated in Figure 4.

4 Concluding Remarks

We see that the effect of hot-start initialization is very significant and suggest that it should be considered as an essential component of every model that contains latent stochastic factors. Specifically to the Heston model, we proposed the two-parameter distribution (3.2), which should be easy to calibrate to market options or to adjust to historical data in the maximum likelihood sense. We wish to stress once more that, even when there is not enough information to calibrate all parameters, the single-value initialization of the latent process is still a poor solution compared to the full-equilibrium hot-start initialization (3.1).

Note that hot-start initialization is also important for the future market scenarios generated by a multifactor model. At first glance, it appears that every scenario is based on some specific simulation path, and thus all the state variables are known, including the latent ones. However, using single-value initialization of the future model will definitely produce an unrealistic implied volatility surface. This is because, in reality, the future market will not know the latent projections of the simulation path. Thus, for a more rigorous generation of scenarios, one has to "forget" the simulated trajectory of the latent variables and to prepare their distribution conditional on the observable trajectory, in the same way as in the maximum-likelihood-related filtrations.

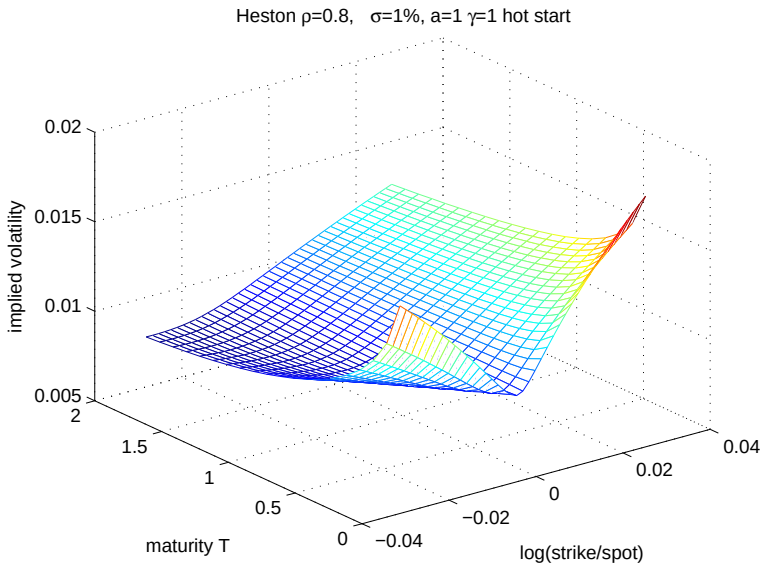

Figure 4: Implied volatility surface for a hot-start Heston model with strong correlation between the stock and the volatility multiplier.

References

[1] Ballestra, L., Ferri, R., and Pacelli, G. The Heston stochastic volatility model for single assets and for asset portfolios: parameter estimation and an application to the Italian financial market. *The International Journal of Business and Finance Research 1*, 2 (2007), 11–23.

[2] Dragulescu, A., and Yakovenko, V. Probability distribution of returns in the Heston model with stochastic volatility. *Quantitative Finance 2* (2002), 443–453.

[3] Forde, M., and Jacquier, A. Small-time asymptotics for implied volatility under the Heston model. *IJTAF 12*, 6 (2009), 861–876.

[4] Gatheral, J. *The Volatility Surface: A Practitioner's Guide*. Wiley Finance, 2006.

[5] Heston, S. A closed-form solution for options with stochastic volatility with application to bond and currency options. *Review of Financial Studies 6* (1993), 327–343.

[6] Hurn, A., Lindsay, K., and McClelland, A. A quasi-maximum likelihood method for estimating the parameters of multivariate diffusions. *Review of Financial Studies 172* (2013), 106–126.

[7] Mechkov, S. Fast-reversion limit of the Heston mode. Available at http://ssrn.com/abstract=2418631, 2014.

Negative Rates: The Challenge and the Opportunity

Ilja Faerman, VP Financial Engineering

Ion Mihai, Quantitative Analyst

NEGATIVE interest rates have recently become an important issue in finance, and pricing methodologies must adjust to this phenomenon. This article discusses the challenges that negative rates pose to the financial community and how Numerix has innovated to address these challenges.

An Overview

On January 15, 2015, the Swiss National Bank (SNB) took the world by surprise by announcing that it was ending the 1.20 floor on the EURCHF exchange rate. (It had put the floor in place in September 2011.) On the same occasion, in a move meant to ease the pressure on the Franc, the SNB announced that it was lowering the interest rate on sight deposits to -0.75% and shifting its target range for the 3M Libor rate downwards by 50 basis points to [-1.25%, -0.25%]. While the SNB move brought negative rates into the headlines, this was not their first appearance. At its previous press conference, the SNB set the interest rate to -0.25% and moved the lower end of its target range for the 3M Libor into negative territory. Since then, the CHF 3M Libor rate has been negative, as shown in Figure 1. Before that, the CHF overnight rate had regularly fixed in the [-0.05%, 0%] range for over three years. Examples of negative rates abound for other currencies as well. For the EUR, the European Central Bank's (ECB) deposit facility rate was set at a negative level since Jun 11, 2014, the overnight EONIA rate and some OIS rates followed suit roughly three months later (when the ECB moved its depo rate further down), and the EURIBOR 1M rate is now negative. This is shown in Figure 2. For the DKK, the Danmarks Nationalbank lowered its interest rate on certificates of deposit to -0.20% in July 2012, and the rate has stayed mostly negative since; the Tomorrow/Next (T/N) rate and the 1M T/N-indexed swap rate have been negative over roughly the same period.

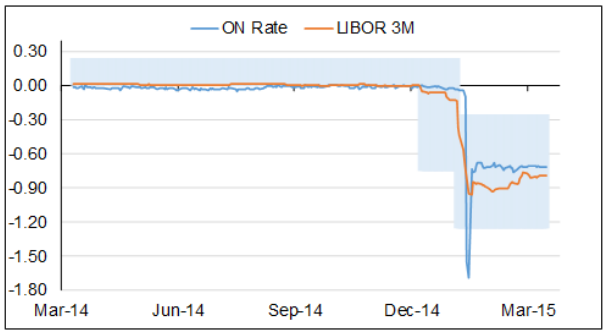

Figure 1: Historical CHF rate between March 2014 and March 2015.

As recently as six months ago, the majority of market participants would have considered the emergence of negative rates a temporary phenomenon. However, the current consensus, especially in light of the SNB's January 15, 2015 announcement, is that rates could stay negative for much longer and could even drop below today's levels.

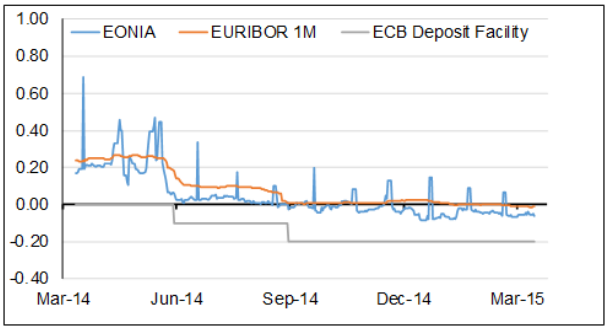

Figure 2: Historical EUR rate between March 2014 and March 2015.

The Challenge

Negative rates impact some of the most basic calculations and procedures used by the financial community. Two prominent examples are the quotation of option volatilities and volatility smile interpolation models. The widespread convention for quoting options is to quote them in terms of implied lognormal volatilities. Underlying the conversion from price to volatility is the assumption that the distribution of the rate is lognormal. It is

a basic mathematical fact that lognormal variables are positive. Therefore, negative rates make the conversion from option price to lognormal volatility impossible. This means that when the market considers it plausible that a rate will become negative, a model which prohibits negative rates is likely to be inappropriate, even if it is only used as a quoting unit conversion tool.

In light of the negative rates environment, the market and the main market data providers have recently embraced two alternatives to the lognormal volatility quotation convention. One is to use normal volatility, which has been around for many years. The assumption here is that the underlying rate follows a normal distribution, which does not restrict the range of the rate to only positive values. It can therefore cope with situations where the forward rate or the strike are negative.

A second alternative, which is becoming widely popular, is to use shifted lognormal volatilities for quoting option prices. Here, instead of making the assumption that the rate is lognormal, one postulates that the rate plus some constant (called the shift) is lognormal. This has the effect of moving the lognormal distribution downwards, into negative territory, by an amount equal to the shift. The shift thus defines the lower bound of the rate: for example, with a shift of 2% the rate can go as low as -2%.

This method is attractive because it is conceptually simple and can be implemented with fairly benign modifications to the existing infrastructure. Nonetheless, shifted lognormal volatilities have some disadvantages. One obvious drawback is that the value of the volatility corresponding to a given option price depends on the chosen shift. The shift value is set somewhat arbitrarily, in accordance with the lowest negative rate perceived possible at that point in time. If the market situation then evolves so that rates become even more negative, then the shift will have to be adjusted upwards. One implication of this is that the values of off-the-grid volatility points, which are obtained through interpolation, will possibly change with the shift (if, as is usually the case, the interpolation is done on volatilities).

Another area where negative rates impact existing standard market practice is smile interpolation models. One of the most widely used of these models is SABR, which stays nonnegative for all admissible values of its β parameter. (In this aspect, the situation is similar to that of the lognormal model.) When the β parameter is less than 0.5, the SABR model can reach 0, but its behavior afterwards is not uniquely defined. This is exploited in the definition of the free-boundary SABR model; see *The Free-Boundary SABR: Natural Extension to Negative Rates* in this issue of Numerix Journal. However, another problem with the SABR model is that, very often, it is implemented through Hagan's formula [2], which is well known to produce arbitrage at low strikes. In an environment where the smile can extend to negative strikes, this issue renders the model unusable.

The Solution

Starting with version 12.0.3, the Numerix CrossAsset analytics library supports the new market practices that have emerged as a response to the negative rates environment. On the market data quotation side, both shifted lognormal volatilities and normal volatilities can now be used as inputs for cap and swaption surfaces and cubes, and the underlying rates and strikes are allowed to be negative. This has also been extended to cap and swaption instruments themselves, which can now consume shifted lognormal or normal volatilities as inputs, or source the volatility from a volatility surface or cube that is defined with such input volatility. These cap and swaption instruments can then be used as calibration instruments for term-structure models or can be priced by an analytic pricer which, according to the type of the input volatility, implements the shifted Black formula or the Bachelier formula.

With the newly added input types, market quotes for option prices can take a variety of forms: not only standard lognormal volatilities, normal volatilities, and the newer shifted lognormal volatilities, but also truncated normal volatilities, premiums, and forward premiums. The interface of each of the Numerix CrossAsset objects that can take one of these market quotes has been revamped to make it easier to use all of the possible

types of market quotes. Specifically, the new interfaces have a `QUOTE` input heading which accepts the market quote. This input heading is always paired with a `QUOTE TYPE` input heading which can be used to specify the type of market quote passed into the `QUOTE` input heading. Depending on the `QUOTE TYPE` input, an additional input may need to be specified to complete the description of the market quote. For instance, with shifted lognormal volatilities, `QUOTE` takes the value of the quoted volatility, `QUOTE TYPE` is set to `SHIFTED LOGNORMAL VOL`, and an additional input, `VOL RATE SHIFT`, is used to hold the value of the shift.

What ⮟ \ Quote Type ➤	NORMAL	LOGNORMAL	SHIFTED LOGNORMAL
Caps/Floors	✓	✓	✓
Swaptions	✓	✓	✓
SABR	✓[1]	✓	✓[2]
AKS SABR	✓[1]	✓	✓[2]
Free-boundary SABR	✓	✓[3]	

Figure 3: Numerix CrossAsset support for various quote types.

1. SABR and AKS SABR work naturally with lognormal volatilities.
2. Shifted versions.
3. Free-boundary SABR works naturally with normal volatilities.

Numerix CrossAsset also introduced new enhancements for the SABR model in the context of the negative rates environment. In a development that goes hand-in-hand with the introduction of shifted lognormal volatilities, shifts are now supported. "Shifted" versions of both the standard SABR model as well as the newer AKS SABR model, which is based on the work of A. Antonov, M. Konikov and M. Spector [1], are available. In addition to the simple extensions with shifts, Numerix CrossAsset now supports the free-boundary SABR model which natively handles negative rates. (This model was developed by the same three Numerix quantitative researchers.) The "shifted" SABR model naturally works with shifted lognormal volatility inputs; the free-boundary SABR model naturally works with normal volatility inputs.

References

[1] Antonov, A., Konikov, M., and Spector, M. SABR spreads its wings. *Risk Magazine* (August 2013), 58–63.

[2] Hagan, P. S., Kumar, D., Lesniewski, A. S., and Woodward, D. E. Managing smile risk. *Wilmott Magazine 1* (November 2002), 84–108.

■

Numerix Real-World Modeling

Ping Sun, Executive Director, Financial Engineering

Vladimir Pavlov, VP Financial Engineering

REAL-WORLD modeling capabilities are essential to many applications in finance, including economic scenario generation and counterparty risk. This article introduces the current status of real-world modeling in Numerix analytics and outlines the roadmap for the future.

The Real-World vs. the Risk-Neutral Measure

Derivatives are valued under the so-called risk-neutral (RN) or pricing measure. Under the RN measure, the values of all assets grow at the same instantaneous rate, the risk-free interest rate. This common-drift property follows from the no-arbitrage assumption, which forms the cornerstone of the modern approach to derivatives valuation.

However, this property does not carry over to real-world asset dynamics. Observable long-term rates of return show considerable variation across assets. Risk premia[1] are generally linked to an asset's exposure to systematic risks and fundamental economic parameters such as investor risk preferences and technology shocks. Risk premia can be positive or negative, depending on the asset's relationship with fundamental risk factors and correlations.

Derivatives pricing models are typically calibrated to be consistent with the observed prices of actively traded products. Derivatives are priced in relative terms. In other words, they are priced in reference to the prices of other assets. This contrasts with the so-called fundamental approach to valuation, wherein assets are valued by using expert and/or quantitative assessments of the economic nature of asset cash-flows and risk. Due to the relative nature of risk-neutral valuation, cross-sectional derivatives pricing data does not contain explicit information about risk premia. Different methods and data sources are therefore needed to extract this information.

The simplest approach to evaluating risk premia is by using long-term average differences in asset returns. These can be complemented by or subjectively adjusted based on expert opinions or fundamental analysis. In the Numerix risk product offering, the results of this type of analysis can be applied by using the projection curves mechanism explained below. The fundamental limitation of this approach is that it does not allow for stochastic risk premia, which could be an important component of the overall market or volatility risk, especially when forecasting over longer horizons.

Current Numerix Real-World Modeling Practices

Current Numerix real-world risk and economic scenario generator (ESG) solutions combine existing Numerix CrossAsset (CA) pricing functionalities with Excel and MATLAB® supplements.

Risk Premia

In the risk-neutral world, expected returns on any investment portfolio are the same (locally) as the risk-free return on the market account. Since no asset can earn an excess return over the risk-free investment, risk premia cannot exist under the RN measure by construction.

Risk premia can be incorporated into projections from CA risk-neutral hybrids by using the existing index curve functionality, which allows for the deterministic shifting of asset value projections. The effective local risk premia are then determined as the difference in forward rates implied by the index versus the model curves.

Model Parameter Estimation

Statistical model calibration is an important aspect of RW functionality—it enables choosing the model which gives the best representation of the

[1] The risk premium is defined as the expected excess return of an asset over the risk-free rate.

observed historical time-series behavior. Prior to CA 12.0, this functionality was not native in CA.

Some econometric estimators have been implemented for specific projects outside of CA. A quasi-maximum likelihood (QML) estimator for the Heston stochastic volatility (SV) model was built in Excel. This estimator relied on a simple approximation of the transition density of the SV model. Another implementation used an even simpler method–of-moments (MOM) estimator.

Bootstrapping the Index Yield Curve

The risk-neutral measure imposes the local expectations theory on bond returns and cannot account for either term premia (excess return on bonds of different maturities) or the time variation in these premia. The user can, however, overwrite model rate projections using CA multi-curve functionality. Specifically, it is possible to add deterministic time-dependent term premia through the projection curve option on rate indices. When applying term premia through a projection curve, users should be mindful of the presence of model-specific convexity, which causes the forward rates implied by the valuation date curve and expected future Libor rates to diverge.

Brute Force (Including Stochastic-on-Stochastic)

Scenario generation is an important part of RW modeling and is usually the first step of stochastic modelling for a wide range of applications in risk and insurance. Generated scenarios are passed into a second modelling layer comprised of revaluation and optimization models. Assuming a simple relationship between the RW model state variables and implied volatilities allows for the conversion of the generated RW scenarios into simulated markets which can then be used to revalue the portfolio and construct the projected distributions of prices and sensitivities.

Real-World Functionality in Numerix CrossAsset 12.0 and Beyond

To capture and extend the basic RW functionality already delivered to clients through custom implementations outside of CA, the following features are currently implemented or planned for CA.

Equity and FX Heston Model with Variance Risk Premium

The Heston model is used extensively in insurance and capital markets. It allows projecting volatility over long time horizons while capturing volatility clustering and generating realistic implied volatility skew dynamics. Defined as the expected equity or foreign exchange return in excess of the risk-free rate, the risk premium is a fundamental quantity in generating economic scenarios.

In the Numerix formulation of the RW Heston model, the risk premium is proportional to the underlying stochastic variance rate, with the variance premium coefficient estimated by a maximum likelihood method (MLE).

Stripping the Index Yield Curve

In order for stochastic short-rate models like Hull–White or CIR to shift future rate projections away from those implied by the risk-free curve at the valuation date, the index curve must be calibrated to match the user input. This is accomplished using global optimization and/or bootstrapping. The final implementation of the bootstrapping method will depend on the results of stability and efficiency tests on the calibration procedure.

RW Hybrid Model

To model multiple underlying assets across currencies in the RW measure, the Numerix hybrid framework has been adapted to RW modeling and scenario generation. The RW hybrid will turn off RN quanto drift adjustments to allow foreign underlyings to grow according to their own risk-free rates and variance risk premia. Path-wise stochastic discount factors reported via Arrow-Debreu price (ADP) indices for each currency within the hybrid framework can be used to check the relevant martingale properties.

Maximum Likelihood Estimation of Heston Model Parameters

The main advantages of econometric estimation for RW modelling are that it guarantees a model providing the best fit to time-series observations and that it allows rigorous hypothesis testing on the model dynamics. Numerix implements QML for the Heston model in the RW measure. The MLE is based on approximating the transition density of the Heston model with the bivariate Gaussian distribution that matches the first and second moments of the time-series.

In its current state, the MLE requires an observable proxy for the volatility. Such proxies may include VIX or other implied volatility-based measures when data on liquid options are available or a time-series of realized volatility constructed from high frequency data. For cases when reliable volatility proxies are not available, a filtering procedure to infer unobserved volatility can be added in future Numerix releases.

Monte Carlo Greeks

An important part of RW modeling is the calculation of future portfolio value and the corresponding sensitivities. The computation of these Greeks usually requires repeatedly bumping the base market and recalculating the value for each simulated scenario, which is very time consuming. To enhance performance, Monte Carlo (MC) Greeks can be calculated on-the-fly and concurrently with the portfolio PV. The current release of CA (12.0) delivers on-the-fly calculations of Delta, Vega, Theta, and Gamma using MC with individual pricing models in EQ and FX. A similar feature under the backward finite-difference method was already available in CA.

Numerix Real-World Roadmap

In addition to working on the planned features, we are also devising a broader roadmap for RW functionality in subsequent Numerix releases, from both a methodology and a modeling perspective. The goal is to build a complete RW modeling framework that will fulfill the full range of requirements for counterparty risk, market risk, and ESG.

Modeling Risk Premia

The first version of the EQ and FX Heston models with variance risk premia was delivered in CA 12.0. This functionality will need to be extended to other EQ and FX models, as well as models for other asset classes including IR, CMDTY, Inflation and Credit. This extension is necessary for real-world modelling over longer projection horizons to capture the excess return on holding a risky underlying.

Modeling Stochastic Volatility Dynamics

Established RN valuation models implemented in CA are quite different from RW market practice, especially when it comes to modeling volatility dynamics. In particular, GARCH models have been widely used in risk applications in the market. Although a large number of GARCH implementations are publically available, none are available in CA.

The following ARCH models deserve consideration:

- GARCH and EGARCH,
- TGARCH, GJR-GARCH: incorporating asymmetric volatility response to positive vs. negative news,
- GARCH–M: incorporating the variance premium in the mean dynamics under GARCH.

Time Series Analysis Tools

MLE for the RW Heston model is available in CA 12.0. CA also provides access to Kalman filter methods in the context of extracting seasonality from historical CMDTY data. In general, a more complete set of time series analysis tools is required, including:

- MLE,
- Generalized Method of Moments (GMM),
- Kalman Filter,
- Copula models of dependence.

More than one method can be combined in a multi-asset cross-currency hybrid in order to achieve more reliable and stable estimation of the model parameters. There are also model-specific features involved in the application of these schemes.

Model Calibration to a User-Specified Mean and/or Distribution of Financial Quantities

CA currently lacks the capability to calibrate a model to match any of several user-specified financial quantities. These user-specified projections may include future yields (delivered in 12.0), expected mean return, percentiles, and other parameters of the return distribution or any other quantities relevant to the user's view of the future state of the market. These projections may cover both price and volatility scenarios.

This functionality would to some extent be similar to the time-series analysis tools. However, there would be more flexibility for users to specify the required projections beyond time series from the historical data.

Some work has been done in this direction in the so-called external calibration functionality. This needs to be developed further and extended in scope.

Consumption of RW Modeling

Generating RW scenarios is just one of the purposes of RW modeling. In the risk space, many quantities are computed under the RW measure, including PFE, VaR, and Expected Shortfall. Ultimately, Numerix will need to build a RW pricing tool in CA to consume the RW models and RW scenarios. The challenge is that some common RN assumptions will need to be revised. The result will be closer to the multicurve framework, which RN pricing is moving towards.

■

99 Park Avenue, 5th Floor
New York, NY 10016
www.numerix.com