

A STOCHASTIC VOLATILITY MODEL FOR RISK-REVERSALS IN FOREIGN EXCHANGE

ALEKSANDAR MIJATOVIĆ AND CLAUDIO ALBANESE

ABSTRACT. It is a widely recognised fact that risk-reversals play a central role in the pricing of derivatives in foreign exchange markets. It is also known that the values of risk-reversals vary stochastically with time. In this paper we introduce a stochastic volatility model with jumps and local volatility, defined on a continuous time lattice, which provides a way of modelling this kind of risk using numerically stable and relatively efficient algorithms.

1. INTRODUCTION

It is generally accepted that in foreign exchange option markets there are at least three main sources of risk. The first is the stochastic behaviour of the FX rate. The Black-Scholes model (Black & Scholes 1973) incorporates this type of uncertainty.

The second source of risk comes from the non-deterministic nature of the volatility of the underlying FX rate. Many attempts have been made to incorporate this observation into the specification of the process driving the FX spot rate. This is usually achieved by extrinsically specifying a diffusion process that governs the behaviour of the volatility for the log-returns of the FX rate. Widely used examples of this approach can be found in (Heston 1993) and (Bates 1996*b*).

The third source of risk in the foreign exchange option markets comes from the observed stochasticity of the skewness of the risk-neutral distribution of the log-returns for the underlying FX spot rate (see (Lipton 2002)). The random behaviour of the skewness for a given currency pair and maturity is reflected in the market price of a 25- Δ risk-reversal¹ (in all that follows we will be referring to the 25- Δ risk-reversal simply as risk-reversal). The problem with modelling this kind of uncertainty is that one cannot simply specify it extrinsically, as is done in the case of stochastic volatility, because that may lead to arbitrage opportunities within the model. In other words, if one wants to capture the stochastic behaviour of the skewness, one has to develop a model that is rich enough to capture this behaviour, is arbitrage-free and numerically tractable. This problem has been tackled by Carr and Wu in (Carr & Wu 2005). They model the jumps up and down in the underlying using two Lévy processes and achieve the stochasticity of the volatility and skewness by randomizing time in both processes.

In this paper we shall describe an alternative approach to the problem. The randomness of volatility and skewness in our model is captured by stochastically changing the local volatility regime in which we are at any moment in time. Our model also includes jumps, which are essential if one wants to reproduce the smiles and skews for short maturities that are observed in the market.

The paper is organized as follows. In section 2 we define the model and outline its properties. Section 3 explains how to price European style derivatives in our framework. Section 4 contains the description of the calibration of the model to three currency pairs: EUR-USD, USD-JPY and

¹For a definition of a 25- Δ risk-reversal see appendix A.

GBP-USD. The options data for these currency pairs (i.e. implied volatilities for the market-defined maturities and strikes) are qualitatively varied in the sense that they are expressing the view of the market on the prices of vanilla options for each currency pair and that these views vary significantly in each of the three cases. The data were deliberately chosen to have this property in order to demonstrate the flexibility of the model. The last section gives the concluding remarks.

2. THE MODEL

Throughout the paper we are assuming that interest rates are given as a deterministic function of calendar time. The domestic (respectively foreign) short rate is denoted by $r^d(t)$ (respectively $r^f(t)$). The FX spot rate X_t will be modelled by a stochastic process with the following components: jumps, local volatility and stochastic switching between volatility regimes which is triggered by X_t hitting specific levels. The model will be defined under the forward measure: the modelled quantity will be the FX forward process given by $F_t = e^{-(r^d(t)-r^f(t))t} X_t$, which is a martingale, rather than the FX spot rate itself.

Our main calibration criterion will be to minimize explicit time-dependence in order to achieve stationarity of the implied volatility surface and avoid the problem of having the surface explicitly related to the FX spot level at which the model was calibrated. This is a well-known problem that affects all pure local volatility models and makes hedge ratios inaccurate. In particular, in the calibration of the model, we will only allow for a minimal non-linear deterministic time change (see subsection (2.4)).

Our model is specified in a largely non-parametric fashion. In order to describe it precisely we will be using the language of functional analysis instead of the usual stochastic calculus. This novel viewpoint, apart from being much simpler, will also give us a vast amount of flexibility when defining the model. In particular it will make transparent the connections between the building blocks of the model and the market features we are trying to capture.

The goal here is to define our model on a continuous time lattice. In order to do that we will make use of spectral theory and numerical linear algebra. These techniques will also come into play in section 3, when we will be faced with the problem of pricing any European payoff. In the rest of this section, we shall go through the steps that are required to build our model on a continuous time lattice.

2.1. The conditional local volatility processes. Our model will consist of m local volatility regimes². A stochastic process, correlated with the level of the forward F_t , will keep track of the volatility state we are in at time t . Corresponding to each volatility state we will have a local volatility process governing the local dynamics of the forward process F_t . We shall return to this point in subsection 2.3 when the stochastic process, mentioned above, will be described in detail. Let us for now assume that at time t we are in the volatility state α , which is one of m possible states.

In this case the forward process F_t can be defined by the following stochastic differential equation

$$(1) \quad dF_t = v_\alpha(F_t)dW_t, \quad \text{where } v_\alpha(F_t) = F_t \min(\sigma_\alpha F_t^{\beta_\alpha - 1}, \bar{\sigma}_\alpha)$$

and W_t is the standard Brownian motion. The constants $\sigma_\alpha, \bar{\sigma}_\alpha$ and β_α need to be specified by calibration for each regime α . The parameter σ_α is closely related to the size of the instantaneous

²In section 4 we use $m = 5$ to calibrate the model to market data for maturities between seven days and five years, but one can, in practice, have many more regimes than that.

variance of the diffusion process in (1), while $\bar{\sigma}_\alpha$ is there to cap the value of the instantaneous log-normal volatility of the process.

It is worth noting that if β_α is equal to one, our conditional local volatility process is the same as the process used for modelling the underlying in the Black-Scholes model. The smile, for this choice of the parameter β_α , is therefore flat. If, on the other hand, β_α is strictly smaller than one, the implied volatility becomes a decreasing function of the strike, which is equivalent to saying that the risk-reversal is positive. If β_α is strictly larger than one the risk-reversal becomes negative, thus causing the implied volatility curve to slope upwards. These considerations are of importance because they make it possible to relate the conditional local volatility behaviour of the model, by choosing the correct value of the parameter β_α , to the view of the market regarding the shape of the smile, conditional on being in the state α .

It is a well-known fact that the probability kernel of diffusion process (i.e. a density function $p(x, t; y, T)$ giving the probability that the value of the process at time T is y , given that, at some prior time t , the value was equal to x) defined by a stochastic differential equation can be described by a corresponding generator. The differential operator which is the generator of the diffusion in (1) takes the form

$$(2) \quad (\mathcal{L}_\alpha u)(F) = \frac{v_\alpha(F)^2}{2} \frac{\partial^2 u}{\partial F^2}(F),$$

for any twice-differentiable function u of F .

Our aim is to build the model on a continuous time lattice. It will soon become clear that the concept of a generator generalizes easily to this kind of situation. In order to build a continuous time lattice we must first discretize the forward rate F_t , which can be achieved as follows. Let Ω be a finite set $\{0, \dots, N\}$ containing the first N integers together with 0 and let $F : \Omega \rightarrow \mathbb{R}$ be a non-negative function which satisfies the following two conditions: $F(0) \geq 0$ and $F(x) > F(x-1)$ for all x in $\Omega - \{0\}$. Given such a function F , the discretized forward rate process F_t can take any of the values $F(x)$, where x is an element in Ω and time t is smaller than some time horizon T .

The next step is to ensure that the dynamics of the discretized forward process correspond to the dynamics specified by the stochastic differential equation (1). This can be achieved by interpreting the generator (2) in the discrete setting. Let the operator $\mathcal{L}_\alpha^\Omega$ be the discretized version of \mathcal{L}_α . The generator \mathcal{L}_α is clearly equal to the Laplace operator multiplied by a scalar function. Since a natural discretization of the Laplace operator is given by

$$(\Delta^\Omega u)(x) = u(x+1) + u(x-1) - 2u(x),$$

for all functions u on Ω , it is clear that the operator $\mathcal{L}_\alpha^\Omega$ can be defined as a tridiagonal matrix of size $(N+1) \times (N+1)$ with entries $\mathcal{L}_\alpha^\Omega(x, y)$, where x and y are elements of Ω . The generator $\mathcal{L}_\alpha^\Omega$ has to satisfy the following conditions for all x in Ω :

$$(3) \quad \sum_{y \in \Omega} \mathcal{L}_\alpha^\Omega(x, y) = 0,$$

$$(4) \quad \sum_{y \in \Omega} \mathcal{L}_\alpha^\Omega(x, y)(F(y) - F(x)) = 0,$$

$$(5) \quad \sum_{y \in \Omega} \mathcal{L}_\alpha^\Omega(x, y)(F(y) - F(x))^2 = v_\alpha(F(x))^2.$$

The first condition is equivalent to the conservation of the probability mass over the infinitesimal time interval dt . The second condition ensures that the forward process is driftless. In other words this condition is equivalent to stipulating that the process F_t is a martingale.

The last condition makes sure that the discretized version of the forward process has the same instantaneous variance as the diffusion defined by the SDE in (1).

Our next task is to obtain the probability kernel of the process F_t from its generator $\mathcal{L}_\alpha^\Omega$. This can be achieved for very general Markov processes by applying spectral methods of operator theory. Here however, we will only illustrate the spectral resolution method in the special case of the operator $\mathcal{L}_\alpha^\Omega$, because this method is sufficient for our purposes and because it can be applied directly to other cases of interest, such as the introduction of jumps (see subsection 2.2). We start by considering the following eigenvalue problem

$$\mathcal{L}_\alpha^\Omega u_n = \lambda_n u_n$$

for the matrix $(\mathcal{L}_\alpha^\Omega(x, y))$. The vectors u_n are the eigenvectors of the linear operator $\mathcal{L}_\alpha^\Omega$ and the scalars λ_n are the corresponding eigenvalues. Except in the trivial cases, the generator $\mathcal{L}_\alpha^\Omega$ will not be a symmetric matrix which implies that the zeros of the characteristic polynomial of $\mathcal{L}_\alpha^\Omega$ (i.e. the eigenvalues λ_n) will not be real. On the other hand it is not hard to see that the eigenvalues of any generator must have a non-positive real part ($\Re(\lambda_n) \leq 0$) and that the complex eigenvalues occur in conjugate pairs (i.e. λ_n is an eigenvalue if and only if $\bar{\lambda}_n$ is an eigenvalue).

In general, of course, there is no guarantee that there exists a complete set of $(N + 1)$ eigenvectors u_n for the operator $\mathcal{L}_\alpha^\Omega$. However, such a set will certainly exist if we can find $(N + 1)$ distinct eigenvalues λ_n of $\mathcal{L}_\alpha^\Omega$. But the set of $(N + 1) \times (N + 1)$ matrices that do not have distinct eigenvalues must have Lebesgue measure zero for the same reason that the set of all polynomials of order $(N + 1)$ with at least two coinciding zeros has Lebesgue measure zero. Therefore we can safely assume that, for a generator specified non-parametrically, the complete set of eigenvectors exists. In the unlikely event that this assumption is not valid, the numerical linear algebra routines needed to solve our lattice model will identify the problem and an arbitrarily small perturbation of a given operator will suffice to rectify the situation. Assuming that there is a solution, the diagonalization problem can be rewritten in the following matrix form

$$\mathcal{L}_\alpha^\Omega = U\Lambda U^{-1},$$

where U is the matrix whose columns are the eigenvectors u_n and Λ is the diagonal matrix with the eigenvalues λ_n .

Key to our constructions is the remark that, if the generator is diagonalizable, we can apply to it an arbitrary function ϕ , defined on the spectrum of the generator, by means of the following formula:

$$(6) \quad \phi(\mathcal{L}_\alpha^\Omega) = U\phi(\Lambda)U^{-1}.$$

Expression (6) is useful because the task of calculating $\phi(\Lambda)$ is a very simple one indeed:

$$\phi(\Lambda) = \begin{pmatrix} \phi(\lambda_0) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \phi(\lambda_N) \end{pmatrix}.$$

The formula in (6) is at the heart of *functional calculus* for generators and plays a pivotal role in our framework for stochastic volatility models.

Formula (6) has many applications. An immediate one allows us to express the probability kernel $p(x, t; y, T) = \mathbb{P}(F_T = F(y) | F_t = F(x))$ of the forward process F_t in the following way

$$p(x, t; y, T) = (e^{(T-t)\mathcal{L}_\alpha^\Omega})(x, y) = \sum_{n=0}^N e^{\lambda_n(T-t)} u_n(x) v_n(y),$$

where the vectors v_n correspond to the columns of the matrix U^{-1} .

2.2. Adding jumps. It is generally accepted that models based on diffusions do not account well for smile effects, especially in the case of short-dated option prices, because of the extremely small likelihoods of large moves in the underlying in short time horizons. In the FX market however, such unexpected jumps are certainly present and their (subjective) probabilities will, to a great extent, shape the prices of short-dated out-of-the-money options. It has therefore been argued (for example in (Bakshi, Cao & Chen 1997)) that the risk-neutral dynamics of the underlying have to be rich enough to include a jump component. In this subsection we shall introduce jumps into our model.

Using spectral theory this can be easily achieved in a general way. What we want is to have different distributions of jump sizes for jumps up and jumps down. Having this property in our model is crucial because the market expectations for jumps up and down are known to the market makers and are almost always very different from each other. Therefore, any process that aspires to model the risk-neutral dynamics of the underlying correctly must be able to account for this difference. The variance gamma model, defined in (Madan, Carr & Chang 1998), has this property since the characteristic function of the underlying process³ is not real.

In general, infinite activity jump processes are built using a special class of stochastic time changes. Such a time change is given by a non-decreasing stationary process T_t with independent increments. The time change T_t is known as a *Bochner subordinator* and is characterized by a *Bernstein function* $\phi(\lambda)$ which has the following property

$$E_0 [e^{-\lambda T_t}] = e^{-\phi(\lambda)t}.$$

For example in the case of the aforementioned variance gamma process, the Bernstein function is of the form

$$(7) \quad \phi(\lambda) = \frac{\mu^2}{\nu} \log \left(1 + \lambda \frac{\nu}{\mu} \right).$$

The parameter μ is the mean-reversion rate and ν is the variance rate of the variance gamma process. In our application the mean-reversion rate μ is normalized to the value 1. Our task now is to do the following: given a generator \mathcal{L} of a Markov process X_t we need to find the generator \mathcal{L}' of the time changed process X_{T_t} . The well-known result from functional analysis (see (Phillips 1952)) tells us that the answer we are looking for is $\mathcal{L}' = -\phi(-\mathcal{L})$.

In our setup we need to add jumps to the conditional local volatility process given by the generator $\mathcal{L}_\alpha^\Omega$. In order to produce asymmetric jumps, we specify two Bernstein functions by choosing two different variance rates in (7): ν_α^+ for jumps up and ν_α^- for jumps down. We then compute separately the two generators

$$\mathcal{L}_\pm = -\phi_\pm(-\mathcal{L}_\alpha^\Omega) = -U_\pm \phi_\pm(-\Lambda) U_\pm^{-1},$$

where Λ is the diagonal matrix from subsection 2.1 and the Bernstein functions ϕ_\pm are given by

$$\phi_\pm(\lambda) = \frac{1}{\nu_\alpha^\pm} \log(1 + \lambda \nu_\alpha^\pm).$$

Each of the square matrices \mathcal{L}_\pm corresponds to a time changed diffusion process. In particular the elements of \mathcal{L}_+ , which are above the diagonal, are the (scaled) probabilities of jumping up in the infinitesimal time interval dt . On the other hand, the sub-diagonal triangle of \mathcal{L}_- contains the (scaled) probabilities of jumping down in the time interval dt . So we can define a new

³The process used in (Madan et al. 1998) is a time-changed Brownian motion with drift. The stochastic time is given by a gamma process.

generator for our process, which will have asymmetric jumps, by combining the two generators in the following way

$$\mathcal{L}_\alpha^\Omega = \begin{pmatrix} d(0,0) & \mathcal{L}_+(0,1) & \cdots & \mathcal{L}_+(0,N-1) & \mathcal{L}_+(0,N) \\ \mathcal{L}_-(1,0) & d(1,1) & \cdots & \mathcal{L}_+(1,N-1) & \mathcal{L}_+(1,N) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathcal{L}_-(N-1,0) & \mathcal{L}_-(N-1,1) & \cdots & d(N-1,N-1) & \mathcal{L}_+(N-1,N) \\ \mathcal{L}_-(N,0) & \mathcal{L}_-(N,1) & \cdots & \mathcal{L}_-(N,N-1) & d(N,N) \end{pmatrix}.$$

Since we want our new process with jumps to be a martingale, we need to make sure that condition (4) is satisfied for the new generator $\mathcal{L}_\alpha^\Omega$. This can be done easily by adjusting the elements just above and below the diagonal of the matrix $\mathcal{L}_\alpha^\Omega$. If, for example, the drift in the x -th row of $\mathcal{L}_\alpha^\Omega$ is positive, we add some probability to the element $\mathcal{L}_\alpha^\Omega(x, x-1)$ so that condition (4) becomes valid. If, on the other hand, the drift in the x -th row is negative, then we can help the generator pull up the process, by adding probability to the element $\mathcal{L}_\alpha^\Omega(x, x+1)$. Once we do this for all x in Ω , the new modified operator, which we again call $\mathcal{L}_\alpha^\Omega$, will satisfy the martingale condition in (4).

The procedure we have just carried out, did not require knowledge of the diagonal elements of $\mathcal{L}_\alpha^\Omega$. They need to be chosen in a way such that the probability conservation (condition (3)) is satisfied. This can be achieved by simply defining the diagonal elements in the following way

$$d(x, x) = - \sum_{y \in \Omega - \{x\}} \mathcal{L}_\alpha^\Omega(x, y).$$

This gives us a well-defined generator $\mathcal{L}_\alpha^\Omega$ for a diffusion process with jumps that can be used to model the risk-neutral forward rate because it is a martingale.

It should be noted that the construction we have just presented differs in two ways from the one in (Madan et al. 1998). The first is that the random time change in our approach can be applied to a general diffusion process that is not necessarily translation invariant⁴. The second is the flexibility in specifying the difference between the distributions of jumps up and jumps down. In our approach we are not confined to a single Bochner subordinator but are free to specify one for each direction of the jump.

2.3. Modelling the dynamics of stochastic volatility. The main idea of our model is to have stochastic volatility that can, not only change the local volatility of the underlying at each moment in time, but specify a jump-diffusion regime as in subsection 2.2 with certain probabilities, which depend on the current volatility level and the current value of the forward rate. We will achieve this in several stages. Let us start by specifying the dynamics of stochastic volatility.

Let V be the set $\{0, \dots, m-1\}$ of all possible volatility states, where m is a positive integer. For each volatility state γ (in V) we define a generator \mathcal{L}_γ^V by specifying the matrix elements $\mathcal{L}_\gamma^V(\alpha, \beta)$, for all $\alpha, \beta \in V$, so that the continuous-time diffusion given by \mathcal{L}_γ^V mean-reverts to the state γ .

We have thus defined m generators \mathcal{L}_γ^V , each of them specifying its own dynamics of the stochastic volatility process. Our next task is to obtain a single global stochastic volatility generator that will favour a certain regime γ conditional on the position of the forward rate F_t . This can be achieved by using a *partition of unity* which is described as follows. Choose a strictly increasing sequence F_γ of the forward rate levels so that, if the forward process F_t is

⁴A key feature of Lévy processes is that they are translation invariant. This is precisely the property of the Brownian motion with drift that is required for the construction in (Madan et al. 1998) to work.

close to the level F_γ , the market views of the smile and skew agree with the ones implied by the process $\mathcal{L}_\gamma^\Omega$ from subsection 2.2. The partition of unity is defined as a sequence of m functions $\epsilon_\gamma : \mathbb{R} \rightarrow [0, 1]$ with the crucial property

$$\sum_{\gamma=0}^{m-1} \epsilon_\gamma(F) = 1 \quad \text{for all } F \in \mathbb{R}.$$

Given the sequence of levels F_γ , such functions can be defined explicitly as piecewise linear functions in the following way:

$$\epsilon_\gamma(F) = \begin{cases} \frac{F-F_{\gamma-1}}{F_\gamma-F_{\gamma-1}} & F \in [F_{\gamma-1}, F_\gamma] \\ \frac{F_{\gamma+1}-F}{F_{\gamma+1}-F_\gamma} & F \in [F_\gamma, F_{\gamma+1}] \\ 0 & \text{otherwise.} \end{cases}$$

This definition has to be modified slightly for the boundary cases when γ equals 0 or $(m-1)$:

$$\epsilon_0(F) = \begin{cases} 1 & F \leq F_0 \\ \frac{F_1-F}{F_1-F_0} & F \in [F_0, F_1] \\ 0 & F \geq F_1, \end{cases} \quad \epsilon_{m-1}(F) = \begin{cases} 0 & F \leq F_{m-2} \\ \frac{F-F_{m-2}}{F_{m-1}-F_{m-2}} & F \in [F_{m-2}, F_{m-1}] \\ 1 & F \geq F_{m-1}. \end{cases}$$

We are now able to define our global generator for the stochastic volatility process, which has the capability of changing its properties when the forward rate undergoes a substantial move. The definition, using the partition of unity we have just defined, is as follows

$$(8) \quad \mathcal{L}_x^V(\alpha, \beta) = \sum_{\gamma=0}^{m-1} \epsilon_\gamma(F(x)) \mathcal{L}_\gamma^V(\alpha, \beta),$$

where α, β are elements in V and $F(x)$ is the function on Ω describing the forward rate process. It follows from the defining property of partition of unity that the matrix \mathcal{L}_x^V is indeed a generator for any element x of the underlying space Ω . As with any stochastic volatility model, our aim is to define a generator \mathcal{L} that will specify the probabilities of going from any state (x, α) in $\Omega \times V$ to any other state (y, β) (of the same set) in the infinitesimal time interval dt . In order to do this we need the *Kronecker delta* function which can be defined on a product of any set with itself, in the following way

$$\delta_{xy} = \begin{cases} 1 & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

We can thus specify the generator \mathcal{L} as

$$\mathcal{L}(x, \alpha; y, \beta) = \mathcal{L}_\alpha^\Omega(x, y) \delta_{\alpha\beta} + \mathcal{L}_x^V(\alpha, \beta) \delta_{xy}.$$

Note that it follows trivially, from the properties of the Kronecker delta, that the matrix \mathcal{L} is a genuine generator. Another important feature is that the generator \mathcal{L} , by definition, does not allow for simultaneous jumps of the state and the volatility variables. This property ensures that our forward process F_t , whose dynamics are specified by \mathcal{L} , remains a martingale.

2.4. Deterministic time-change. The model we have described so far is completely stationary, i.e. its implied volatility surface has no explicit time dependence. Because of this feature it is possible to calibrate it to the market prices (i.e. the implied volatilities of the 25- Δ calls and puts and the at-the-money European options for all market-specified maturities from seven days to five years) only with accuracy of ± 5 basis points. In other words, the volatilities implied by the model differ from the prices observed in the market by at most 0.05%.

In order to get a better match for the 25- Δ and the at-the-money implied volatilities we have to introduce a minimal deterministic time change. This amounts to specifying an increasing

function $f : [0, T] \rightarrow [0, \infty)$ that will deterministically transform calendar time t to financial time $f(t)$ (in the examples of section 4 the time horizon T equals five years).

Let t be a tenor⁵ date (i.e. a market-specified date between seven days and five years) and let σ_+, σ_- and σ_0 be the market-implied volatilities for the 25- Δ call, 25- Δ put and the at-the-money call respectively, all maturing at t . Assume further that σ_u, σ_d and σ_a are the Black-Scholes volatilities implied by our model for the same set of options. We now define the transformation f , at the tenor date t , to be

$$f(t) = \frac{\sigma_+^2 + \sigma_-^2 + \sigma_0^2}{\sigma_u^2 + \sigma_d^2 + \sigma_a^2} t.$$

Using this formula we can define the function f for all tenor dates and then extend linearly to the entire interval between now and five years. In section 3 we will see that the deterministic time change we have just described enters into the probability kernel of our model in a very isolated and controlled way (see equation (9)).

Since our model, before the time change is introduced, already captures well the features of the underlying market, the function f we have just defined will only differ slightly from the identity. This can be seen best from figures 5, 10 and 15, which contain graphs of the function f that was required to calibrate the model to each of the currency pairs in section 4.

3. PRICING

In order to tackle the pricing problem for European payoffs in our framework, we need to calculate the transition probabilities $p((x, \alpha), t; (y, \beta), T)$, where $(x, \alpha), (y, \beta) \in \Omega \times V$ and T is a time horizon, implied by the generator \mathcal{L} which was defined in subsection 2.3.

Let $U_s = (u((x, \alpha), s; (y, \beta), S))$ be a stochastic matrix induced by the generator \mathcal{L} . The coordinate functions $u((x, \alpha), s; (y, \beta), S)$ are the transition probabilities for the underlying process starting at the state (x, α) and ending at the state (y, β) in the time interval $[s, S]$. The quantity s denotes the financial time and can therefore be expressed as $s = f(t)$, where f is the function from subsection 2.4 and t is the calendar time. Similarly $S = f(T)$ is the financial time horizon. It is well-known that in this case U_s satisfies the backward Kolmogorov equation

$$\frac{\partial U_s}{\partial s} + \mathcal{L}U_s = 0$$

with the final condition $U_S = \mathbb{I}$ (\mathbb{I} is the identity matrix on the vector space \mathbb{R}^k where $k = m(N + 1)$). The solution of the backward Kolmogorov equation can therefore be expressed in terms of functional calculus as $U_s = \exp((S - s)\mathcal{L})$ and can be calculated explicitly using the spectral decomposition of the operator \mathcal{L} . The first step is to calculate the eigenvalues λ_n of the generator \mathcal{L} and the second is to find the eigenvectors u_n of \mathcal{L} , collect them into a matrix U , and find the columns of the inverse U^{-1} which we will denote by v_n . In the examples of section 4 the matrix \mathcal{L} has a dimension of $m(N + 1) = 355$. For matrices of this size, diagonalization routines such as `dgeev` in LAPACK are very efficient.

Once we have the spectral decomposition of \mathcal{L} , we can calculate the original probability kernel, which depends on calendar time, using functional calculus in the following way

$$(9) \quad p((x, \alpha), t; (y, \beta), T) = e^{(f(T) - f(t))\mathcal{L}}((x, \alpha), (y, \beta)) = \sum_{n=1}^{m(N+1)} e^{\lambda_n(f(T) - f(t))} u_n(x, \alpha) v_n(y, \beta).$$

It is not hard to see that eigenvalues λ_n have a negative real part. Therefore in the case of long dated options only very few eigenvalues will play a role, because the exponential of a negative

⁵All tenor dates we used are listed in the legend of figure 2.

number becomes negligibly small very quickly. Another important fact which follows from (9) is that the pricing kernel depends in a very isolated way on the (financial) time to maturity.

The price C_t of a European payoff $h(X_T)$, where X_t is a value of an exchange rate at time t and T is the time maturity, can be calculated in the following way

$$C_t = e^{-(r^d(T)T - r^d(t)t)} \sum_{(y, \beta) \in \Omega \times V} p((x, \alpha), t; (y, \beta), T) h(e^{-(r^d(T) - r^f(T))T} F(y)).$$

The point x from Ω is chosen so that $e^{-(r^d(t) - r^f(t))t} F(x) = X_t$ and α in V corresponds to the volatility regime in which we are at time t .

4. CALIBRATION

We shall calibrate the model for three different currency pairs: EUR-USD, USD-JPY and GBP-USD. For a fixed maturity, the most liquid instruments in the FX markets are the at-the-money call, the 25- Δ put and the 25- Δ call (see for example (Lipton 2002)). For each of the currency pairs and all market-specified maturities (as listed in the legends of figures 2, 7 and 12) we use the market-implied Black-Scholes volatilities as a benchmark to set the parameters of our model in such a way as to reprice all the above options correctly. Our main calibration criterion is to minimize the explicit time dependence in order to preserve the correct (i.e. market-implied) smile and skew through time. Our model is flexible enough to capture the risk-neutral dynamics of each of the three currency pairs even though they are qualitatively different from each other. The stationarity of the calibration makes sure that the forward smile and skew, even after a large move of the underlying FX rate, still have the desired shape.

In order to calibrate our model we select a non-homogenous grid with $N + 1 = 71$ points, used to span the possible values of the forward rate F_t . We also choose $m = 5$ local volatility regimes in order to capture the correct behaviour of the risk-reversals of all three currency pairs. We shall now describe the calibration procedure for each currency pair.

4.1. EUR-USD. A cursory inspection of the market data for EUR-USD, plotted in figure 1, reveals that for short maturities the risk-reversal is negative and wanes gradually as the time to maturity increases, disappearing completely at the five year mark. This tells us that the skewness of the market-implied risk-neutral distribution of the underlying is present for shorter maturities and that this distribution becomes more symmetric with time. Since the risk-reversals are positive, the left tail of the implied pdf must be fatter than that on the right. It should also be noted that the price of the 25- Δ call appreciates, relative to the value of the at-the-money option, with increasing time to maturity. This is a consequence of the fact that the vega risk for out-of-the-money options increases with time.

Using this intuition about the risk-neutral dynamics of the underlying we can proceed to the calibration. The following table specifies the values of the model parameters that have been found to work best for EUR-USD.

α	σ_α	β_α	$\bar{\sigma}_\alpha$	ν_α^-	ν_α^+	F_α
0	9.300%	-100%	500%	0.140%	0.000%	1.17
1	8.900%	5%	500%	0.070%	0.000%	1.20
2	8.700%	100%	500%	0.000%	0.000%	1.25
3	8.825%	190%	500%	0.000%	0.055%	1.32
4	9.100%	280%	500%	0.000%	0.090%	1.38

These parameters specify the generators $\mathcal{L}_\alpha(x, y)$ for the local volatility processes with jumps as defined in subsection 2.2. They also describe the partition of unity functions $\epsilon_\gamma(F)$, where $\gamma \in V = \{0, \dots, 4\}$, from subsection 2.3. In order to define the global generators $\mathcal{L}_x^V(\alpha, \beta)$ for stochastic volatility dynamics given by equation (8), we need to specify the local mean-reverting generators $\mathcal{L}_\gamma^V(\alpha, \beta)$ that govern the dynamics of stochastic volatility conditional upon the volatility process being in the regime γ . These generators are given in the following table.

$$\begin{aligned} \mathcal{L}_0^V &= \begin{pmatrix} -13 & 13 & 0 & 0 & 0 \\ 14 & -27 & 13 & 0 & 0 \\ 0 & 14 & -23 & 9 & 0 \\ 0 & 0 & 14 & -23 & 9 \\ 0 & 0 & 0 & 14 & -14 \end{pmatrix}, & \mathcal{L}_1^V &= \begin{pmatrix} -14 & 14 & 0 & 0 & 0 \\ 9 & -22 & 13 & 0 & 0 \\ 0 & 14 & -23 & 9 & 0 \\ 0 & 0 & 14 & -23 & 9 \\ 0 & 0 & 0 & 14 & -14 \end{pmatrix}, \\ \mathcal{L}_2^V &= \begin{pmatrix} -14 & 14 & 0 & 0 & 0 \\ 9 & -23 & 14 & 0 & 0 \\ 0 & 9 & -18 & 9 & 0 \\ 0 & 0 & 14 & -23 & 9 \\ 0 & 0 & 0 & 14 & -14 \end{pmatrix}, & \mathcal{L}_3^V &= \begin{pmatrix} -14 & 14 & 0 & 0 & 0 \\ 9 & -23 & 14 & 0 & 0 \\ 0 & 9 & -23 & 14 & 0 \\ 0 & 0 & 13 & -22 & 9 \\ 0 & 0 & 0 & 14 & -14 \end{pmatrix}, \\ & & \mathcal{L}_4^V &= \begin{pmatrix} -14 & 14 & 0 & 0 & 0 \\ 9 & -23 & 14 & 0 & 0 \\ 0 & 9 & -23 & 14 & 0 \\ 0 & 0 & 13 & -27 & 14 \\ 0 & 0 & 0 & 13 & -13 \end{pmatrix} \end{aligned}$$

TABLE 1. Local generators $\mathcal{L}_\gamma^V(\alpha, \beta)$ for stochastic volatility used to calibrate the model to the options data for EUR-USD.

All the parameters in the model are chosen in a way such that there is no explicit time dependence in their calculation and yet all the known option prices can be reobtained from the model. This requires good understanding of the market data and of the market-implied risk-neutral dynamics, an approach that is quite different from what is usually proposed in the standard literature on calibration (e.g. (Lipton 2002)).

Using the market-implied Black-Scholes volatilities we have identified the forward levels F_α where the likelihood of the regime change is maximal. Each of the regimes has its own local volatility and jump dynamics. The EUR-USD exchange rate, at the time when the option data were recorded, equaled 1.2193. Our starting regime is regime 0. The negative β_0 is there to account for the relatively large positive risk-reversal. Since a diffusion process alone cannot support such a skew within a very short time horizon, we need non-zero jump-down intensity ν_α^- which will give us the required shape of the implied volatility curve.

If, for example, a large move up to level 1.25 happens, the most likely local volatility regime we will be in will have a flat smile. This of course does not mean that our model would then yield a flat implied volatility curve, because the risk-neutral dynamics of the model would still be implicitly influenced by other regimes.

Another thing to notice is that the implied probability distribution functions (see figure 4) are well behaved in the sense that they have a single maximum⁶ and hence the risk-neutral dynamics of the underlying are governed purely by the thickness and asymmetry of their tails.

⁶If one attempts to simulate smile and skew effects by taking a convex combination of different models, one is invariably faced with the problem of “bumpy” probability distribution functions with more than one maximum.

4.2. **USD-JPY.** It is clear, just by glancing at the market data in figure 6, that this currency pair has a large persistent positive risk-reversal that increases with time to maturity. The values of the parameters β_α are negative in order to keep the skews of all the local volatility regimes decreasing in strike. The values of the parameters are as follows.

α	σ_α	β_α	$\bar{\sigma}_\alpha$	ν_α^-	ν_α^+	F_α
0	7.600%	-335%	500%	0.000%	0.000%	92.00
1	8.000%	-60%	500%	0.000%	0.000%	97.00
2	8.800%	-335%	500%	0.500%	0.500%	101.25
3	8.200%	-180%	500%	0.900%	0.200%	107.32
4	10.500%	-40%	500%	0.900%	0.000%	110.00

As in subsection 4.1, we need to specify the numerical values for the generators \mathcal{L}_γ^V that govern the stochastic volatility dynamics. They are given by the matrices in table 2.

$$\mathcal{L}_0^V = \begin{pmatrix} -9 & 9 & 0 & 0 & 0 \\ 19.6 & -28.6 & 9 & 0 & 0 \\ 0 & 19.6 & -24.6 & 5 & 0 \\ 0 & 0 & 19.6 & -24.6 & 5 \\ 0 & 0 & 0 & 19.6 & -19.6 \end{pmatrix}, \mathcal{L}_1^V = \begin{pmatrix} -19.6 & 19.6 & 0 & 0 & 0 \\ 5 & -14 & 9 & 0 & 0 \\ 0 & 19.6 & -24.6 & 5 & 0 \\ 0 & 0 & 19.6 & -24.6 & 5 \\ 0 & 0 & 0 & 19.6 & -19.6 \end{pmatrix},$$

$$\mathcal{L}_2^V = \begin{pmatrix} -19.6 & 19.6 & 0 & 0 & 0 \\ 5 & -24.6 & 19.6 & 0 & 0 \\ 0 & 5 & -10 & 5 & 0 \\ 0 & 0 & 19.6 & -24.6 & 5 \\ 0 & 0 & 0 & 19.6 & -19.6 \end{pmatrix}, \mathcal{L}_3^V = \begin{pmatrix} -19.6 & 19.6 & 0 & 0 & 0 \\ 5 & -24.6 & 19.6 & 0 & 0 \\ 0 & 5 & -24.6 & 19.6 & 0 \\ 0 & 0 & 9 & -14 & 5 \\ 0 & 0 & 0 & 19.6 & -19.6 \end{pmatrix},$$

$$\mathcal{L}_4^V = \begin{pmatrix} -19.6 & 19.6 & 0 & 0 & 0 \\ 5 & -24.6 & 19.6 & 0 & 0 \\ 0 & 5 & -24.6 & 19.6 & 0 \\ 0 & 0 & 9 & -28.6 & 19.6 \\ 0 & 0 & 0 & 9 & -9 \end{pmatrix}$$

TABLE 2. Local generators $\mathcal{L}_\gamma^V(\alpha, \beta)$ for stochastic volatility used to calibrate the model to the options data for USD-JPY.

The prevailing FX rate for USD-JPY, when the market data were recorded, was equal to 110.415. Each of the five regimes is centred around a level F_α . Since the jumps are required to sustain the skew for short maturities, we only need non-zero values for jump-intensities ν_α^- and ν_α^+ for the regimes with centres F_α close to the above exchange rate.

4.3. **GBP-USD.** This currency pair has relatively symmetric market-implied risk-neutral dynamics with a positive risk-reversal that does not change much with time to maturity. The parameters required to capture these dynamics are as follows.

α	σ_α	β_α	$\bar{\sigma}_\alpha$	ν_α^-	ν_α^+	F_α
0	9.700%	230%	500%	0.700%	0.000%	1.67
1	9.600%	60%	500%	0.450%	0.000%	1.73
2	8.900%	-100%	500%	0.100%	0.000%	1.79
3	9.600%	20%	500%	0.000%	0.150%	1.89
4	10.900%	170%	500%	0.000%	0.600%	2.00

The stochastic volatility generators \mathcal{L}_γ^V used in the calibration of the model to the implied volatility surface for GBP-USD are the same as the ones for EUR-USD and are given in table 1.

5. CONCLUSIONS

In this paper we have presented a new approach to modelling the behaviour of risk-reversals in foreign exchange. Our model is a stochastic volatility model with jumps and local volatility, which captures the stochasticity of risk-reversals by having the ability to move continuously from one local volatility regime to another. The richness of the model allows us to calibrate it to the market and retain a nearly stationary behaviour. The numerical solubility of the model is based on the theory of continuous-time Markov chains, which lends itself well to the application of numerical linear algebra.

Another interesting feature of the model is that it gives reasonable arbitrage-free prices for vanilla options that are struck very far out of the money (the implied volatilities are never above 35%). The reason for this lies in the fact that the risk-neutral density implied by the model is uniformly bounded above and below by two log-normal probability density functions that correspond to the values of the implied volatilities that can be observed in the markets.

APPENDIX A. BASIC MARKET CONVENTIONS IN FX

In foreign exchange markets vanilla options are referred to in terms of their Black-Scholes deltas and their prices are quoted in terms of the implied volatility. So, a 25- Δ call with an implied volatility of 10% is an out-of-the-money call option with a strike K such that, when we calculate the Black-Scholes delta at K using the implied volatility of 10%, we find that its value is 0.25. This method of referring to vanilla options in terms of their deltas and quoting their prices in terms of their Black-Scholes volatilities yields a transparent and user-friendly way of comparing option prices across currency pairs.

The at-the-money call (or put) with the at-the-money volatility σ_0 , expiring at time t , is an option struck at K such that its Black-Scholes delta, when evaluated at K and σ_0 , equals $\frac{1}{2}e^{-r^f(t)t}$. The value of a 25- Δ risk-reversal⁷ is just the difference $\sigma_p - \sigma_c$, where σ_p is the current market value of the 25- Δ put and σ_c is the market value of the 25- Δ call. For more information on the foreign exchange markets see (Lipton 2001).

REFERENCES

- Bakshi, G., C. Cao & Z. Chen (1997), ‘Empirical performance of alternative option pricing models’, *Journal of finance* **52**, 2003–2049.
- Bates, D. (1996a), ‘Dollar jump fears, 1984-1992: Distributional abnormalities implicit in foreign currency futures options’, *Journal of International Money and Finance* **15**, 65–93.
- Bates, D. (1996b), ‘Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options’, *Review of Financial Studies* **9**, 69–107.
- Black, F. & M. Scholes (1973), ‘The pricing of options and corporate liabilities’, *Journal of Political Economy* **81**, 637–654.
- Carr, P. & L. Wu (2005), ‘Stochastic skew in currency options’, *Working Paper*.

⁷In this paper we will be referring to this structure simply as risk-reversal.

- Heston, S. (1993), 'Closed-form solution for options with stochastic volatility, with application to bond and currency options', *Review of Financial Studies* **6**, 327–343.
- Lipton, A. (2001), *Mathematical Methods for Foreign Exchange*, World Scientific Publishing.
- Lipton, A. (2002), 'The vol smile problem', *Risk February*, 61–65.
- Madan, D., P. Carr & E.C. Chang (1998), 'The variance gamma process and option pricing', *European Finance Review* **2**(1), 79–105.
- Phillips, R.S. (1952), 'On the generation of semigroups of linear operators', *Pacific Journal of Mathematics* **2**(3), 343–369.

ALEKSANDAR MIJATOVIĆ, DEPARTMENT OF MATHEMATICS, IMPERIAL COLLEGE LONDON.
E-mail address: a.mijatovic@imperial.ac.uk

CLAUDIO ALBANESE, INDEPENDENT CONSULTANT




FIGURE 1. Term structure of implied volatilities for EUR-USD at-the-money options, 25- Δ puts and 25- Δ calls. The triangles \blacktriangle represent market data for the specific dates, while the continuous curves graph the implied volatility of the model as a function of time.

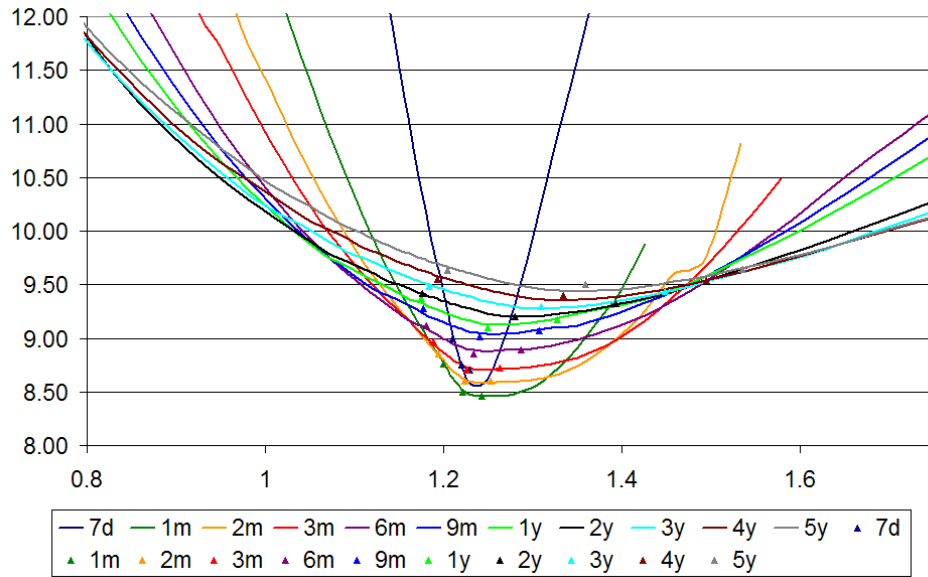


FIGURE 2. Extrapolated implied volatilities for EUR-USD European options for generic strikes. For each market specified maturity the triangles \blacktriangle represent market-implied volatilities for the following strikes: the 25- Δ call, the 25- Δ put and the at-the-money call or put. The continuous curves graph the implied volatility of the model as a function of strike.




FIGURE 3. Implied volatilities for extreme strikes for maturities between one year and five years in EUR-USD.




FIGURE 4. EUR-USD probability density function under the forward measure.




FIGURE 5. Deterministic time change $f(t)$ (measured in years) as a function of calendar time t (also in years) for EUR-USD.




FIGURE 6. Term structure of implied volatilities for USD-JPY at-the-money options, 25- Δ puts and 25- Δ calls. The triangles \blacktriangle represent market data for the specific dates, while the continuous curves graph the implied volatility of the model as a function of time.




FIGURE 7. Extrapolated implied volatilities for JPY-USD European options for generic strikes. For each market-specified maturity the triangles \blacktriangle represent market-implied volatilities for the following strikes: the 25- Δ call, the 25- Δ put and the at-the-money call or put. The continuous curves graph the implied volatility of the model as a function of strike.




FIGURE 8. Implied volatilities for extreme strikes for maturities between nine months and five years in USD-JPY.




FIGURE 9. USD-JPY probability density function under the forward measure.




FIGURE 10. Deterministic time change $f(t)$ (measured in years) as a function of calendar time t (also in years) for USD-JPY.




FIGURE 11. Term structure of implied volatilities for GBP-USD at-the-money options, 25- Δ puts and 25- Δ calls. The triangles \blacktriangle represent market data for the specific dates, while the continuous curves graph the implied volatility of the model as a function of time.

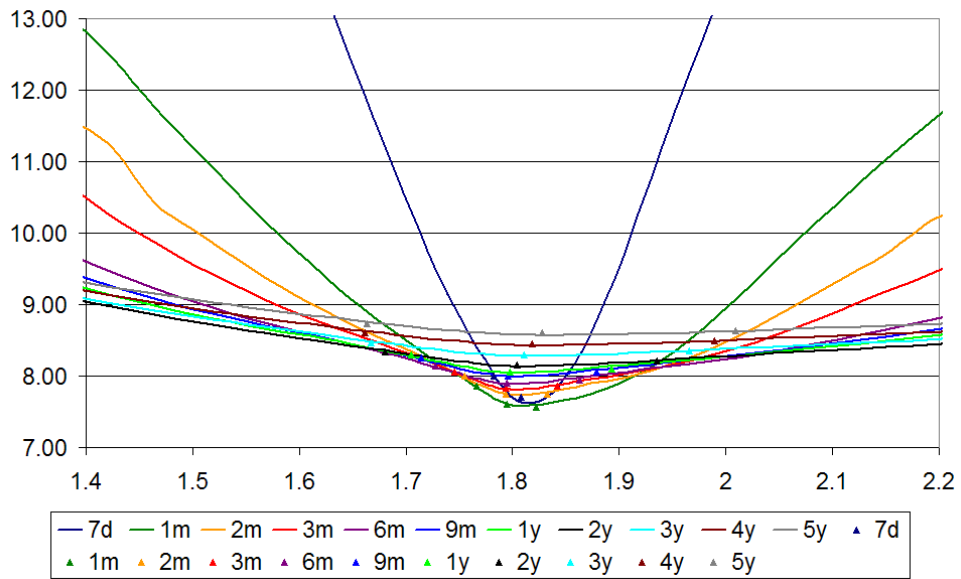


FIGURE 12. Extrapolated implied volatilities for GBP-USD European options for generic strikes. For each market-specified maturity the triangles \blacktriangle represent market-implied volatilities for the following strikes: the 25- Δ call, the 25- Δ put and the at-the-money call or put. The continuous curves graph the implied volatility of the model as a function of strike.




FIGURE 13. Implied volatilities for extreme strikes for maturities between one year and five years in GBP-USD.





FIGURE 14. GBP-USD probability density function under the forward measure.

FIGURE 15. Deterministic time change $f(t)$ (measured in years) as a function of calendar time t (also in years) for GBP-USD.