Fast Ninomiya-Victoir calibration of the Double-Mean-Reverting Model

Jim Gatheral
(joint work with Christian Bayer and Morten Karlsmark)

Baruch COLLEGE
The City University of New York

Stochastic Processes and their Statistics in Finance
Okinawa, Japan
October 31, 2013

Overview of this talk

- Specification of the DMR model
- Calibration of model parameters
- The Ninomiya-Victoir scheme and extensions
- Simulation of the DMR model
- Examples of model fits to options data
- Summary and conclusions

Modeling SPX and VIX

- It is well-known that the empirically observed implied volatility surface is not consistent with Black-Scholes.
- Many models have been proposed proposed to fit market implied volatilities better and describe the dynamics of the volatility surface.
  - For example, local volatility models, Lévy models, stochastic volatility models, stochastic volatility models with jumps and so on.
- With the advent of trading in VIX options in 2006 however, marginal risk-neutral densities of forward volatilities of SPX became effectively observable, substantially constraining possible choices of volatility dynamics.
- Various authors have since proposed models that price both options on SPX and options on VIX more or less consistently with the market.

The DMR model

- In [my Bachelier 2008 presentation], a specific three factor variance curve model was introduced with dynamics motivated by economic intuition for the empirical dynamics of the variance.
- In this *double-mean-reverting* or *DMR* model, the dynamics are given by

$$\mathrm{d}S_t = \sqrt{v_t} S_t \mathrm{d}W_t^1, \tag{1a}$$
$$\mathrm{d}v_t = \kappa_1 (v_t' - v_t)\, \mathrm{d}t + \xi_1\, v_t^{\alpha_1}\, \mathrm{d}W_t^2, \tag{1b}$$
$$\mathrm{d}v_t' = \kappa_2 (\theta - v_t')\, \mathrm{d}t + \xi_2\, v_t'^{\alpha_2}\, \mathrm{d}W_t^3, \tag{1c}$$

where the Brownian motions W_i are all in general correlated with $\mathbb{E}[\mathrm{d}W_t^i\, \mathrm{d}W_t^j] = \rho_{ij}\, \mathrm{d}t$.

Qualitative features of the DMR model

- Instantaneous variance v mean-reverts to a level v' that itself moves slowly over time with the state of the economy, mean-reverting to the long-term mean level θ.
- Also, it is a stylized fact that the distribution of volatility (whether realized or implied) should be roughly lognormal
  - When the model is calibrated to market option prices, we find that indeed $\alpha_1 \approx 1$ consistent with this stylized fact.
- As we will see later, the DMR model calibrated jointly to SPX and VIX options markets fits pretty well.

Computations in the DMR model

- One drawback of the DMR model is that calibration is not easy
  - No closed-form solution for European options exists so finite difference or Monte Carlo methods need to be used to price options.
  - Calibration is therefore slow.
- In [my Bachelier 2008 presentation], the DMR model is calibrated using an Euler-Maruyama Monte Carlo scheme with the partial truncation step of [Lord, Koekkoek, and van Dijk].
- In this talk, we show how to apply the Monte Carlo scheme of [Ninomiya and Victoir] to the calibration of the DMR model, substantially improving calibration time.

Model calibration

- The DMR model has many parameters:
  - One could argue that it is both mis-specified and over-parameterized.
- In [my Bachelier 2008 presentation], the parameters of the DMR model were calibrated to the VIX and SPX options markets with a sequence of steps that we will now individually describe.

Variance swaps from the log-strip

- Under diffusion assumptions, the fair value of a variance swap is given by evaluating the so-called log-strip of European puts and calls (see Chapter 11 of [The Volatility Surface] for example):

$$\mathbb{E}\left[\int_t^T v_s \, \mathrm{d}s \middle| \mathcal{F}_t\right] = 2\left\{\int_{-\infty}^0 p(k)\, dk + \int_0^\infty c(k)\, dk\right\}, \tag{2}$$

where $k = \log(K/F_{t,T})$ is the log-strike and p and c respectively are put and call prices expressed as a fraction of the strike price.

- Variance swaps may thus be estimated from historical option prices by interpolation, extrapolation and integration.

Estimation of κ_1, κ_2, θ and ρ_{23}

- In the DMR model, the fair strike of a variance swap is given by the expression

$$\mathbb{E}\left[\int_t^T v_s\,\mathrm{d}s \middle| \mathcal{F}_t\right] = \theta\,\tau + (v_t - \theta)\frac{1 - e^{-\kappa_1\,\tau}}{\kappa_1} + (v'_t - \theta)\frac{\kappa_1}{\kappa_1 - \kappa_2}\left\{\frac{1 - e^{-\kappa_2\,\tau}}{\kappa_2} - \frac{1 - e^{-\kappa_1\,\tau}}{\kappa_1}\right\} \tag{3}$$

  which is affine in the state variables v_t and v'_t.
- Fixing θ, κ_1 and κ_2, and given daily variance swap estimates, time series of v_t and v'_t may be imputed by linear regression.
  - Optimal values of θ, κ_1 and κ_2 are obtained by minimizing mean squared differences between the fitted and actual variance swap curves.

Calibrated parameters

- With daily data from January 2001 to April 2008, the optimal choice of parameters was found to be

$$\begin{aligned} \theta &= 0.078, \\ \kappa_1 &= 5.5, \\ \kappa_2 &= 0.10. \end{aligned}$$

- The correlation ρ_{23} between W_t^2 and W_t^3 was then estimated as the historical correlation between the series v_t and v_t'. The estimated value was

$$\rho_{23} = 0.59.$$

Motivation for fitting SABR

- It seems that volatility dynamics are roughly lognormal
  - Option prices and time series analysis lead us to the same conclusion.
- The SABR model of [Hagan, Kumar, Lesniewski, and Woodward] is the simplest possible lognormal stochastic volatility model
  - And there is an accurate closed-form approximation to implied volatility.
- The lognormal SABR process is:

$$\begin{aligned} \frac{dS}{S} &= \Sigma\, dZ \\ \frac{d\Sigma}{\Sigma} &= \nu\, dW \end{aligned} \tag{4}$$

  with $\langle dZ, dW \rangle = \rho\, dT$.
- Fitting SABR might allow us to impute effective parameters for a more complicated model.

The SABR formula

As shown originally by Hagan et al., to lowest order in time to expiration, the solution to (4) in terms of the Black-Scholes implied volatility σ_{BS} is approximated by:

$$\sigma_{BS}(k) = \sigma_0\, f\left(\frac{k}{\sigma_0}\right)$$

where $k := \log(K/F)$ is the log-strike and

$$f(y) = -\frac{\nu\, y}{\log\left(\frac{\sqrt{\nu^2\, y^2 + 2\,\rho\,\nu\, y + 1} - \nu\, y - \rho}{1-\rho}\right)}$$

It turns out that this simple formula is reasonably accurate for longer expirations too.

- Note that the formula is independent of time to expiration T.

The term structure of ν

As of 25-Apr-2008, plot fitted ν for each slice against T_{exp}:

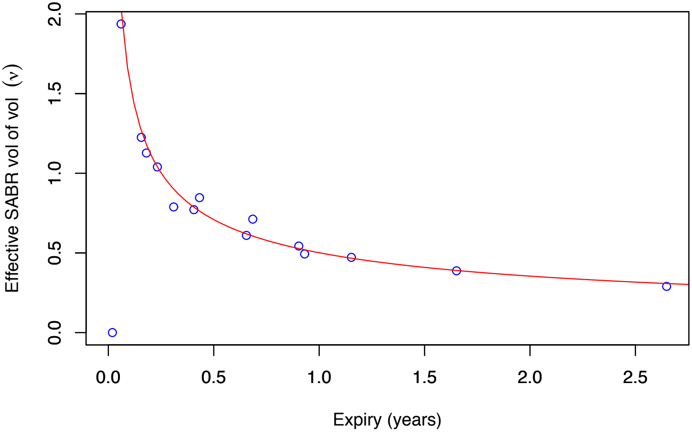

The red line is the function $\frac{0.501}{\sqrt{T}}$.

Estimation of the exponents α_1 and α_2

- The exponents α_1 and α_2 control how volatility of volatility changes with the volatility level.
- To obtain a proxy for the volatility of volatility we note that the lognormal SABR model (with $\beta = 1$) tends to fit the smile at any given expiration very well.
  - One of the SABR parameters is the volatility of volatility ν.
- We note further that empirically, the term structure of ν is given by

$$\nu(\tau) \approx \frac{\nu_{\mathit{eff}}}{\sqrt{\tau}}$$

  - ν_{eff} may then be used as a proxy for volatility of volatility.
- We use VIX as a proxy for the level of volatility.

SABR fits to SPX: ν_{eff}

Figure 1: Computing ν_{eff} every day for seven years gives the following time-series plot:

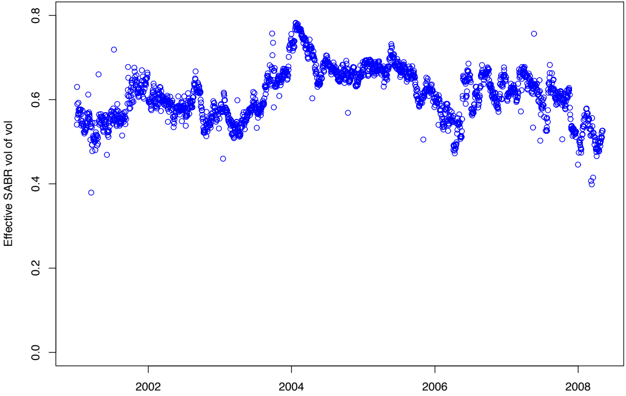

Observations from ν_{eff} time-series

- Lognormal volatility of volatility ν_{eff} is empirically rather stable
  - The dynamics of the volatility surface imply that volatility is roughly lognormal.
- Can we see any patterns in the plot?
  - For example, does ν_{eff} depend on the level of volatility?

Regression of ν_{eff} vs VIX

Figure 2: Regression does show a pattern!

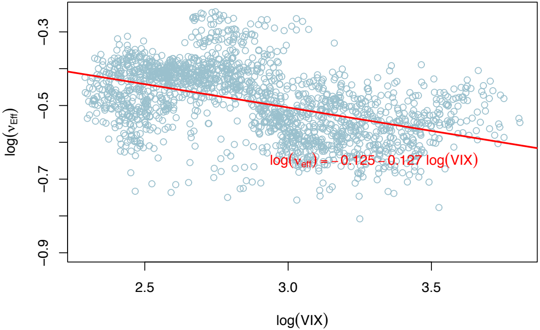

- $VIX \sim \sqrt{v}$ so we conclude that $\alpha_1 \approx 0.94$ – not so far from lognormal!
- We don't have enough data to pin down α_2 so we set $\alpha_2 \approx 0.94$.

Daily calibration of remaining parameters

- Although the volatility of volatility parameters ξ_1 and ξ_2 are in principle constants of the DMR model, it is clear from Figure 1 that they are not constant in the data.
  - And of course, VIX option prices fluctuate from day to day.
  - The implied volatility of VIX options is like volatility of volatility.
- This leads to the following daily procedure:
  1. Calibrate v_t and v'_t to variance swaps (from SPX option prices) using linear regression.
  2. Calibrate ξ_1 and ξ_2 to VIX options data.
  3. Calibrate the correlations ρ_{12} and ρ_{13} to the SPX volatility surface.
- In steps 2 and 3, we need to use numerical techniques to compute VIX and SPX options prices respectively.

The LKV scheme

- In [my Bachelier 2008 presentation], calibration of ξ_1, ξ_2, ρ_{12} and ρ_{13} was performed using Monte-Carlo simulation.
- The chosen scheme was Euler-Maruyama with partial truncation as in [Lord, Koekkoek, and van Dijk]:

$$
\begin{aligned}
x((k+1)\Delta) &= -\frac{1}{2}v(k\Delta)\Delta + \sqrt{v(k\Delta)}Z_k^1 \\
\tilde{v}((k+1)\Delta) &= \tilde{v}(k\Delta) + \kappa_2\left(\tilde{v}'(k\Delta) - \tilde{v}(k\Delta)\right)\Delta + \left(\tilde{v}(k\Delta)^+\right)^{\alpha_1} Z_k^2 \\
\tilde{v}'((k+1)\Delta) &= \tilde{v}'(k\Delta) + \kappa_2\left(\theta - \tilde{v}'(k\Delta)\right)\Delta + \left(\tilde{v}'(k\Delta)^+\right)^{\alpha_2} Z_k^3.
\end{aligned}
$$

- In the above, Δ is the time step, $v(k\Delta) = \tilde{v}(k\Delta)^+$, $v'(k\Delta) = \tilde{v}'(k\Delta)^+$, $x(k\Delta) = \log(S(k\Delta))$, $Z_k^i \sim N(0,\Delta)$ and $\mathbb{E}[Z_k^i Z_k^j] = \rho_{ij}\Delta$.
- Calibration of the DMR model with this scheme is slow!

The Ninomiya-Victoir scheme

- In [Ninomiya and Victoir], a general second order weak discretization scheme for stochastic differential equations was introduced.
- Consider a multi-dimensional stochastic differential equation in Stratonovich form

$$\mathrm{d}\mathbf{X}(t,\mathbf{x}) = V_0(\mathbf{X}(t,\mathbf{x}))\mathrm{d}t + \sum_{i=1}^{d} V_i(\mathbf{X}(t,\mathbf{x})) \circ \mathrm{d}B_t^i, \tag{5}$$

where $\mathbf{X}(0,\mathbf{x}) = \mathbf{x} \in \mathbb{R}^N$, $B_t^1, \ldots, B_t^d$ are d *independent* standard Brownian motions and $V_i : \mathbb{R}^N \to \mathbb{R}^N$, $i = 0, \ldots, d$, are sufficiently regular vector fields.

The Ninomiya-Victoir scheme

- The Ninomiya-Victoir scheme is given by

$$
\begin{aligned}
&\mathbf{X}^{(NV)}(0,\mathbf{x}) = \mathbf{x},\\
&\mathbf{X}^{(NV)}((k+1)\Delta,\mathbf{x})\\
&= \begin{cases} e^{\frac{\Delta}{2}V_0} e^{Z_k^1 V_1} \cdots e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}\left(k\Delta,\mathbf{x}\right), & \Lambda_k = -1,\\ e^{\frac{\Delta}{2}V_0} e^{Z_k^d V_d} \cdots e^{Z_k^1 V_1} e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}\left(k\Delta,\mathbf{x}\right), & \Lambda_k = +1. \end{cases}
\end{aligned} \tag{6}
$$

- $e^{tV}\mathbf{x} \in \mathbb{R}^N$ denotes the solution at time $t \in \mathbb{R}$ to the ODE

$$\dot{\mathbf{y}} = V\left(\mathbf{y}\right), \quad \mathbf{y}\left(0\right) = \mathbf{x},$$

  i.e. the *flow* of the vector field.
- The Λ_k take values ± 1 with probability $1/2$, the Z_k^j are independent $\mathcal{N}\left(0,\Delta\right)$ random variables.

Features of the Ninomiya-Victoir scheme

- One step in the NV scheme corresponds to a (non-discrete) cubature formula of order $m = 5$ in the sense of [Lyons and Victoir].
- One can also interpret the NV scheme as the stochastic version of a classical operator splitting scheme, where the infinitesimal generator $\mathcal{L} = V_0 + \frac{1}{2}\sum_{i=1}^{d} V_i^2$ of the diffusion is split into the first order differential operator V_0 and the second order differential operators $\frac{1}{2}V_1^2, \ldots, \frac{1}{2}V_d^2$.
- The NV scheme is now widely used in applications such as Inria's software PREMIA for financial option computations.
  - In particular, a variant of the NV scheme due to [Alfonsi] can be used to simulate the Heston model.

Solving the ODEs

- Cubature methods, and the Ninomiya-Victoir scheme in particular, need us to solve ODEs quickly and accurately.
- General cubature methods involve ODEs with a rather complicated structure, involving all vector-fields at all times.
  - Numerical ODE methods such as Runge-Kutta are typically required.
- If we are very lucky, all of the ODE flows may be solved exactly – in terms of easy-to-evaluate expressions.
- In such a case, one has effectively found a second order weak approximation method which can be implemented without relying on numerical ODE solvers
  - The Ninomiya-Victoir method can be expected to perform especially well in such cases.

The Drift Trick

- The Heston model is one of the lucky cases where all of the ODEs may be solved in closed form.
- However, one soon encounters models where the ODEs have no closed-form solution
  - For example, the SABR model.
- [Bayer, Friz, and Loeffen] observed that the class of favorable models can be significantly enlarged by working with an almost trivial modification of the NV scheme.

The idea of the Drift Trick

- We can rewrite the SDE as

$$\mathrm{d}\mathbf{X}(t,\mathbf{x}) = \left(V_0(\mathbf{X}(t,\mathbf{x})) - \sum_{j=1}^{d} \gamma_j V_j(\mathbf{X}(t,\mathbf{x})) \right) \mathrm{d}t$$

$$+ \sum_{j=1}^{d} V_j(\mathbf{X}(t,\mathbf{x})) \circ \mathrm{d}\left(B_t^j + \gamma_j t \right)$$

$$=: V_0^{(\gamma)}(\mathbf{X}(t,\mathbf{x}))\,\mathrm{d}t + \sum_{j=1}^{d} V_j(\mathbf{X}(t,\mathbf{x})) \circ \mathrm{d}\left(B_t^j + \gamma_j t \right)$$

  whatever the choice of drift parameters $\gamma_1, \ldots, \gamma_d$.
- In many cases of interest, it's possible to choose the γ_i so as to permit the solution of all ODEs in closed-form.
  - In particular, the DMR model with $\alpha_1 = \alpha_2 = 1$, the *Double Lognormal* model.

The NV scheme with drift trick

- The Ninomiya-Victoir scheme with drift trick is now given by

$$
\begin{aligned}
\mathbf{X}^{NVd)}(0,\mathbf{x}) &= \mathbf{x}, \\
\mathbf{X}^{(NVd)}((k+1)\Delta,\mathbf{x}) &= \\
&\begin{cases}
e^{\frac{\Delta}{2}V_0^{(\gamma)}} e^{Z_k^1 V_1} \dots e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_0^{(\gamma)}} \mathbf{X}^{(NVd)}(k\Delta,\mathbf{x}), & \Lambda_k = -1, \\
e^{\frac{\Delta}{2}V_0^{(\gamma)}} e^{Z_k^d V_d} \dots e^{Z_k^1 V_1} e^{\frac{\Delta}{2}V_0^{(\gamma)}} \mathbf{X}^{(NVd)}(k\Delta,\mathbf{x}), & \Lambda_k = +1,
\end{cases}
\end{aligned}
\tag{7}
$$

  where the $Z_k^i \sim \mathcal{N}(\Delta\gamma_i, \Delta)$ are again independent of each other.
- This amended scheme corresponds to splitting $\mathcal{L}$ according to

$$
\mathcal{L} = V_0 + \frac{1}{2}\sum_{i=1}^{d} V_i^2 = V_0^{(\gamma)} + \sum_{i=1}^{d}\left\{\frac{1}{2}V_i^2 + \gamma_i V_i\right\}.
$$

More operator splitting

- In models such as the DMR model with $\alpha_1, \alpha_2 \neq 1$, the drift trick is not enough to permit closed-form solution of the drift ODE.
- We may then try to find vector fields $V_{0,1}$ and $V_{0,2}$ such that $V_0 = V_{0,1} + V_{0,2}$ and the ODEs driven by $V_{0,1}$ and $V_{0,2}$ have (closed-form) solutions $e^{tV_{0,1}}$ and $e^{tV_{0,2}}$, respectively.
- In that case, the solution $e^{\Delta V_0}$ of the ODE driven by the vector field V_0 at time Δ can be approximated by

$$e^{\Delta V_0}\mathbf{x} = e^{\Delta V_{0,2}} e^{\Delta V_{0,1}}\mathbf{x} + \mathcal{O}(\Delta^2),$$

  a method sometimes known as the *symplectic Euler scheme*.
- One contribution of the present work is to show that the NV scheme can be further extended in this way whilst maintaining second order weak convergence.

Our modified Ninomiya Victoir (NVs) scheme

- In our modification of the NV scheme, applying the symplectic Euler scheme to the solution of the drift ODE, we iterate according to

$$\mathbf{X}^{(NVs)}((k+1)\Delta, \mathbf{x}) =$$
$$\begin{cases} e^{\frac{\Delta}{2}V_{0,1}} e^{\frac{\Delta}{2}V_{0,2}} e^{Z_k^1 V_1} \dots e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_{0,2}} e^{\frac{\Delta}{2}V_{0,1}} \mathbf{X}^{(NVs)}(k\Delta, \mathbf{x}), & \Lambda_k = -1, \\ e^{\frac{\Delta}{2}V_{0,1}} e^{\frac{\Delta}{2}V_{0,2}} e^{Z_k^d V_d} \dots e^{Z_k^1 V_1} e^{\frac{\Delta}{2}V_{0,2}} e^{\frac{\Delta}{2}V_{0,1}} \mathbf{X}^{(NVs)}(k\Delta, \mathbf{x}), & \Lambda_k = +1. \end{cases} \tag{8}$$

- This modified NVs scheme again has second order convergence in the weak sense.

- In the case of the DMR model with $\alpha_1, \alpha_2 \neq 1$, this further splitting of the drift operator V_0 is sufficient to permit us to solve all of the ODEs in closed form.

Itô formulation of the DMR model

The DMR model (1) re-expressed in terms of independent Brownian motions B^i reads:

$$
\begin{aligned}
\mathrm{d}S_t &= \sqrt{v_t}\, S_t\, \mathrm{d}B_t^1, \\
\mathrm{d}v_t &= \kappa_1 \left(v_t' - v_t\right) \mathrm{d}t + \xi_1 v_t^{\alpha_1} \left(\tilde{\rho}_{1,2}\mathrm{d}B_t^1 + \sqrt{1-\tilde{\rho}_{1,2}^2}\,\mathrm{d}B_t^2\right), \\
\mathrm{d}v_t' &= \kappa_2 \left(\theta - v_t'\right) \mathrm{d}t + \xi_2 v_t'^{\alpha_2} \left(\tilde{\rho}_{1,3}\mathrm{d}B_t^1 + \tilde{\rho}_{2,3}\mathrm{d}B_t^2 + \sqrt{1-\tilde{\rho}_{1,3}^2-\tilde{\rho}_{2,3}^2}\,\mathrm{d}B_t^3\right),
\end{aligned} \tag{9}
$$

where $\tilde{\rho}_{12} = \rho_{12}$, $\tilde{\rho}_{13} = \rho_{13}$ and $\tilde{\rho}_{23} = \frac{\rho_{23}-\rho_{12}\rho_{13}}{\sqrt{1-\rho_{12}^2}}$.

Stratonovich formulation of the DMR model

- To apply NVs, we need the Stratonovich formulation:

$$\mathbf{X}(t,\mathbf{x}) = \mathbf{x} + \int_0^t V_0\left(\mathbf{X}(s,\mathbf{x})\right) ds + \sum_{j=1}^{3} \int_0^t V_j\left(\mathbf{X}(s,\mathbf{x})\right) \circ dB_s^j \tag{10}$$

  where the state vector $\mathbf{X}(t,\mathbf{x}) = (S_t, v_t, v_t')^T$, and the initial condition is $\mathbf{x} = (S_0, v_0, v_0')^T$.
- The driving vector fields $\{V_0, V_1, V_2, V_3\}$ are given explicitly in the following slide.

Explicit expressions for the vector fields

We have

$$V_0(\mathbf{x}) = \begin{pmatrix} -\frac{1}{2}\left(\frac{1}{2}\,\xi_1\,\tilde{\rho}_{1,2}\,x_2^{\alpha_1-\frac{1}{2}}\,x_1 + x_2\,x_1\right) \\ -\kappa_1\,(x_2 - x_3) - \frac{1}{2}\,\xi_1^2\,\alpha_1\,x_2^{2\alpha_1-1} \\ -\kappa_2(x_3-\theta) - \frac{1}{2}\,\xi_2^2\,\alpha_2\,x_3^{2\alpha_2-1} \end{pmatrix}$$

and also

$$V_1(\mathbf{x}) = \begin{pmatrix}\sqrt{x_2}\,x_1 & \tilde{\rho}_{1,2}\,\xi_1\,x_2^{\alpha_1} & \tilde{\rho}_{1,3}\,\xi_2\,x_3^{\alpha_2}\end{pmatrix}^T$$

$$V_2(\mathbf{x}) = \begin{pmatrix}0 & \sqrt{1-\tilde{\rho}_{1,2}^2}\,\xi_1\,x_2^{\alpha_1} & \tilde{\rho}_{2,3}\,\xi_2\,x_3^{\alpha_2}\end{pmatrix}^T$$

$$V_3(\mathbf{x}) = \begin{pmatrix}0 & 0 & \sqrt{1-\tilde{\rho}_{1,3}^2-\tilde{\rho}_{2,3}^2}\,\xi_2\,x_3^{\alpha_2}\end{pmatrix}^T.$$

Solving the ODEs

- In order to implement the NVs scheme, we thus need to solve the ODEs

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{x}(t) = V_i(\mathbf{x}(t))$$

for all $i \in \{0,\ 1,\ 2,\ 3\}$ and $t \in \mathbb{R}$ with some given boundary condition.

- It is relatively straightforward to solve the ODEs for $i \in \{1, 2, 3\}$ in closed form.
- Solving the ODE

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{x}(t) = V_0(\mathbf{x}(t))$$

requires further splitting.

The flow of the Stratonovich drift vector field

- To solve the ODE for $i = 0$, we write

$$V_0 = V_{0,1} + V_{0,2}$$

with

$$V_{0,1}(x) = \begin{pmatrix} -\frac{1}{2}\, x_2\, x_1 \\ -\kappa_1 \left(x_2 - x_3\right) \\ -\kappa_2 \left(x_3 - \theta\right) \end{pmatrix},$$

$$V_{0,2}(x) = \begin{pmatrix} -\frac{1}{4}\, \xi_1\, \tilde{\rho}_{1,2}\, x_2^{\alpha_1 - \frac{1}{2}}\, x_1 \\ -\frac{1}{2}\, \xi_1^2\, \alpha_1\, x_2^{2\alpha_1 - 1} \\ -\frac{1}{2}\, \xi_2^2\, \alpha_2\, x_3^{2\alpha_2 - 1} \end{pmatrix}.$$

- It is again straightforward to solve the corresponding ODEs in closed form.

The double lognormal case: $\alpha_1 = \alpha_2 = 1$

- In the Double Lognormal case with $\alpha_1 = 1$, $\alpha_2 = 1$, we may applying the drift trick to get closed-form ODE solutions.
- Specifically, with

$$V_0^\gamma = V_0 - \gamma_1 V_1 - \gamma_2 V_2 - \gamma_3 V_3,$$

and choosing

$$\gamma_1 = -\xi_1 \tilde{\rho}_{1,2},$$

$$\gamma_2 = -\frac{\kappa_1 + \frac{1}{2}\xi_1^2 + \gamma_1 \tilde{\rho}_{1,2} \xi_1}{\xi_1 \sqrt{1 - \tilde{\rho}_{1,2}^2}},$$

$$\gamma_3 = -\frac{\kappa_2 + \frac{1}{2} \xi_2^2 - \tilde{\rho}_{1,3} \xi_2 \gamma_1 - \tilde{\rho}_{2,3} \xi_2 \gamma_2}{\xi_2 \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2}},$$

we end up with much simpler expressions for the vector fields.

The adjusted drift vector field

- Explicitly, after applying the drift trick, we get

$$V_0^{\gamma} = \begin{pmatrix} -\frac{1}{2}\, x_2\, x_1 \\ \kappa_1\, x_3 \\ \kappa_2\, \theta \end{pmatrix}$$

which is much simpler than the original

$$V_0(\mathbf{x}) = \begin{pmatrix} -\frac{1}{2}\left(\frac{1}{2}\,\xi_1\,\tilde{\rho}_{1,2}\,\sqrt{x_2}\,x_1 + x_2\,x_1\right) \\ -\kappa_1\,(x_2 - x_3) - \frac{1}{2}\,\xi_1^2\,x_2 \\ -\kappa_2(x_3 - \theta) - \frac{1}{2}\,\xi_2^2\,x_3 \end{pmatrix}$$

- It is again straightforward to compute solutions to these ODEs in closed-form.

Daily model fitting

- Once again, the model parameters κ_1, κ_2, θ and ρ_{23} are considered fixed.
- The state variables v_t and v'_t are obtained by linear regression against the fair values of variance swaps proxied by the log-strip.
  - Arbitrage-free interpolation and extrapolation of the volatility surface is achieved using the SVI parameterization in [Gatheral and Jacquier].
- The volatility-of-volatility parameters ξ_1 and ξ_2 are obtained by calibrating the DMR model to the market prices of VIX options (using NVs).
- The correlation parameters ρ_{12} and ρ_{13} are then calibrated to SPX options.

Pricing VIX options

- The payoff of a call option on the VIX index with strike K expiring at time T may be written as

$$\left(\sqrt{\mathbb{E}\left[\int_T^{T+\Delta} v_s \mathrm{d}s \,\middle|\, \mathcal{F}_T\right]} - K\right)^+$$

  where Δ is is roughly one month.
- Each Monte Carlo path generates a value for v_T and v'_T, so the expected forward variance $\mathbb{E}\left[\int_T^{T+\Delta} v_s \mathrm{d}s \,\middle|\, \mathcal{F}_T\right]$ is given by equation (3).
  - Averaging over all paths gives the model price of the VIX option.

Calibration to VIX options

- Our chosen objective function is the sum of squared differences between market VIX implied volatilities and model VIX implied volatilities. Errors are weighted by the reciprocal of the bid-ask spread:

$$\sqrt{\sum_i \left(\frac{\sigma_i^{mid} - \sigma_i^{model}}{\sigma_i^{ask} - \sigma_i^{bid}} \right)^2}.$$

Calibration to SPX options

- We are then left with the correlation parameters ρ_{12} and ρ_{13} to calibrate to the SPX volatility surface.
  - Note only two parameters to fit the entire volatility surface!
- Our objective function is again the sum of squared differences between market SPX implied volatilities and model SPX implied volatilities, weighted by the reciprocal of the bid-ask spread.

Two days in history

- We pick two days in history to fit the DMR model, one before the 2008 financial crisis, and one after:
  - April 3, 2007 and September 15, 2011.
- Recall that fixed model parameters were as follows:

θ	0.078
κ_1	5.5
κ_2	0.10
ρ_{23}	0.59
α_1	0.94
α_2	0.94

Fitted parameters

- With v_t, v_t' from variance swaps, ξ_1, ξ_2 from VIX options, and ρ_{12}, ρ_{13} from SPX options we obtain:

	03-Apr-2007	15-Sep-2011
v	0.0153	0.114
v'	0.0224	0.110
ξ_1	2.873	2.689
ξ_2	0.302	0.502
ρ_{12}	-0.992	-0.982
ρ_{13}	-0.615	-0.727

- Note that fitted parameters from two very different market environments are very similar.

Variance swap fit as of April 3, 2007

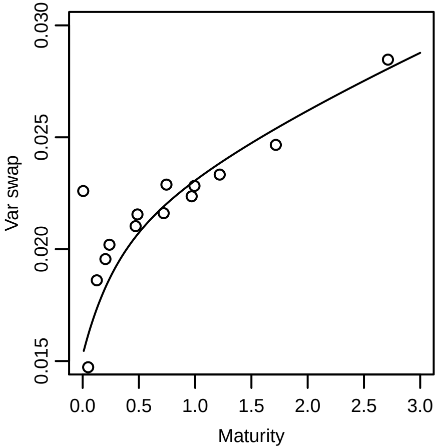

Figure 3: The points are SPX variance swaps (from the log-strip), the solid curve is the DMR model fit.

VIX fit as of April 3, 2007

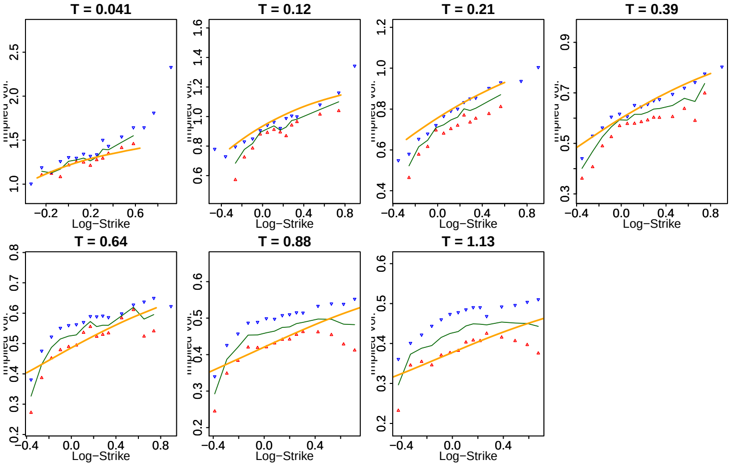

Figure 4: VIX smiles as of April 3, 2007: Bid vols in red, ask vols in blue, mid vols in green, and model fits in orange.

SPX fit as of April 3, 2007

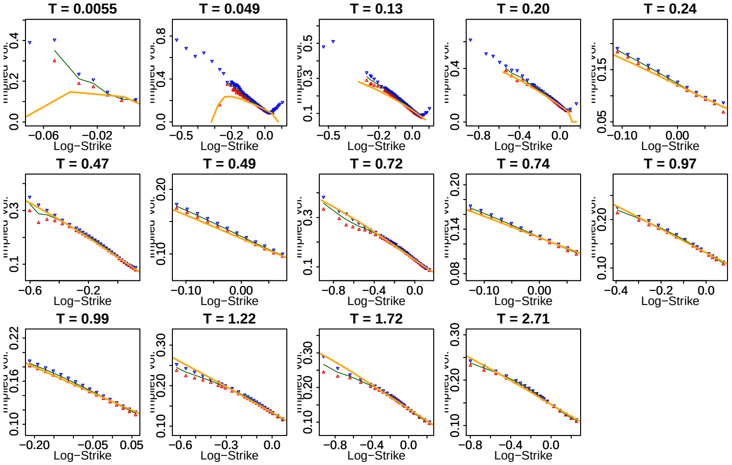

VIX fit as of September 15, 2011

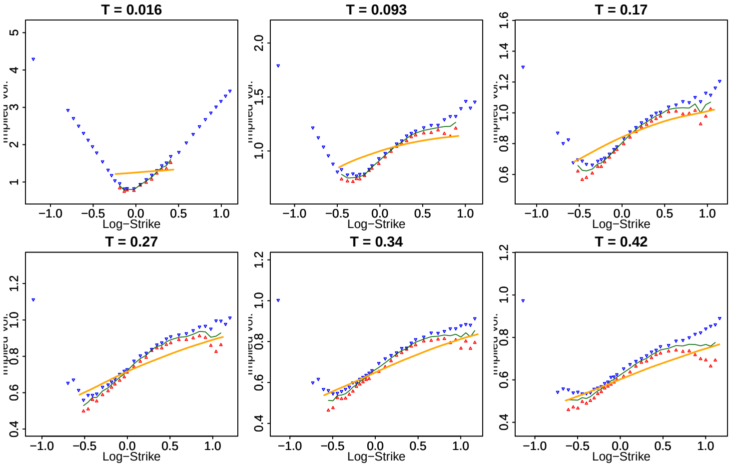

Figure 5: VIX smiles as of September 15, 2011: Bid vols in red, ask vols in blue, and model fits in orange. Note that the fitted smiles seem a little too flat.

SPX fit as of September 15, 2011

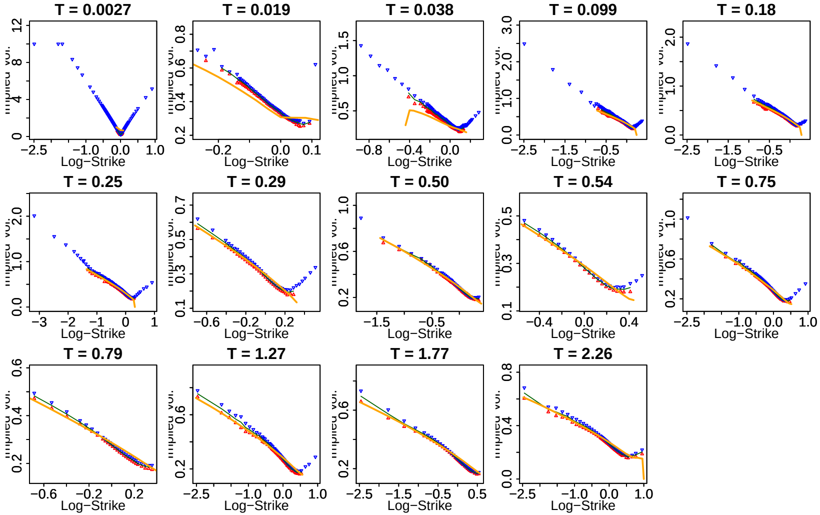

Double Lognormal model calibration to 2011 data

- In Figure , we saw that the DMR model with $\alpha_1 = \alpha_2 = 0.94$ generates VIX option smiles that are too flat.
- This motivates us to calibrate the simpler Double Lognormal version of the DMR model with $\alpha_1 = \alpha_2 = 1$.
  - Simulation is also faster because we use the drift trick. Each time step is less complex.

VIX fit of Double Lognormal as of September 15, 2011

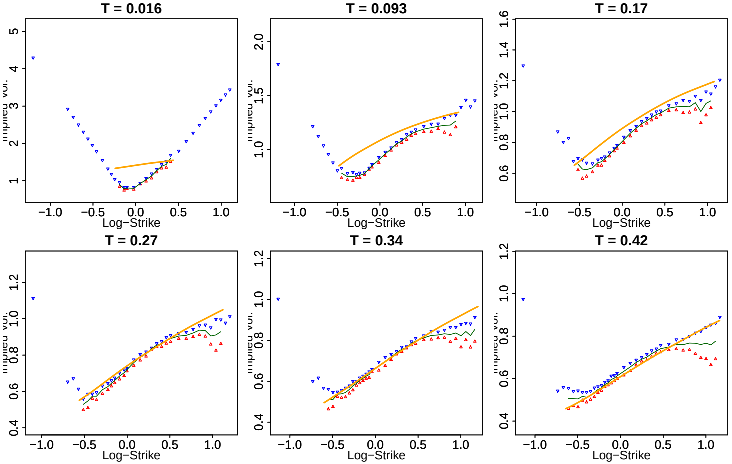

Figure 6: VIX smiles as of September 15, 2011: Bid vols in red, ask vols in blue, and Double Lognormal model fits in orange.

Remarks on the Double Lognormal fit

- The smiles got a little steeper.
- The algorithm (with the drift trick) is less complex.
- And the model with $\alpha_1 = \alpha_2 = 1$ is more parsimonious.

- Double Lognormal seems like the better choice!

Performance tradeoff between NV and EM

- The Ninomiya-Victoir (NV) scheme permits us to achieve a given target RMSE with fewer time steps than Euler-Maruyama (EM).
- However, the computational cost of each NV time step is greater than EM.
- The tradeoff between NV and EM must therefore be assessed experimentally.

	2D	3D
$\alpha_1 = \alpha_2 = 0.94$	4.55	6.84
$\alpha_1 = \alpha_2 = 1$	1.81	3.08

Table 1: Relative computation times for NV steps in terms of EM steps. 2D means simulation of the variance process only (*i.e.* for VIX options); 3D means simulation of the full model. The values are obtained by simulating with 90 time steps and 2^{18} QMC paths using the parameters obtained in the 2011 calibrations.

Optimal calibration recipe

- From experiment, we conclude that it is better to use the EM discretization when calibrating to SPX options where there is little if any RMSE reduction benefit from using the NV step.
- However, for VIX options, we can achieve a speedup of $3-4$ times in the 2007 example, 2 in the 2011 example and 5 in the 2011 lognormal DMR example.
- The optimal calibration recipe appears to be:
    - Calibrate ξ_1 and ξ_2 with a Ninomiya-Victoir scheme.
    - Calibrate ρ_{12} and ρ_{13} with an Euler-Maruyama scheme.
- Using Java code with 30 time steps and 2^{11} paths we can typically calibrate the model to both SPX and VIX option markets in less than 5 seconds.

Conclusion

- We have presented two straightforward modifications of the standard Ninomiya-Victoir discretization scheme that conserve second order weak convergence but permit simple closed-form solutions to the ODE's.
  - NV with drift trick, and NV with extra splitting of the drift vector field.
- Using these schemes for VIX options and the simpler Euler-Maruyama scheme for SPX options, we demonstrated that it is possible to achieve fast and accurate calibration of the DMR model to both SPX and VIX options markets simultaneously.
- Moreover, we demonstrated that the DMR model fits SPX and VIX options market data well for two particular dates chosen to represent two very different market environments from before and after the 2008 financial crisis.
  - The fitted parameters of the model over time appear to be remarkably stable.

References

Aurélien Alfonsi, High order discretization schemes for the CIR process: Application to affine term structure and Heston models, *Mathematics of Computation* **79**(269) 209–237 (2010).

Christian Bayer, Peter Friz, and Ronnie Loeffen, Semi-closed form cubature and applications to financial diffusion models, *Quantitative Finance* **13**(5), 769–782 (2013).

Christian Bayer, Jim Gatheral, and Morten Karlsmark, Fast Ninomiya-Victoir calibration of the double-mean-reverting model, *Quantitative Finance* forthcoming (2013).

Jim Gatheral, *The Volatility Surface: A Practitioner's Guide*, John Wiley and Sons, Hoboken, NJ (2006).

Jim Gatheral, *Consistent Modeling of SPX and VIX Options*, Fifth World Congress of the Bachelier Finance Society (2008).

Jim Gatheral and Antoine Jacquier, Arbitrage-free SVI volatility surfaces, *Quantitative Finance* forthcoming (2013).

Patrick Hagan, Deep Kumar, Andrew Lesniewski, and Diana Woodward, Managing smile risk, *Wilmott Magazine* **1**/(1) 84–108 (2002).

Roger Lord, Koekkoek, and D. van Dijk, A comparison of biased simulation schemes for stochastic volatility models *Quantitative Finance* **10**(2) 177–194 (2010).

Terry Lyons and Nicolas Victoir, Cubature on Wiener space, *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **460**(2041) 169–198 (2004).

Syoiti Ninomiya and Nicolas Victoir, Weak approximation of stochastic differential equations and application to derivative pricing, *Applied Mathematical Finance* **15** 107–121 (2008).