

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

Christian Bayer*, Jim Gatheral†, Morten Karlsen‡

April 26, 2013

Abstract

We consider the three factor double mean reverting (DMR) model of Gatheral (2008), a model which can be successfully calibrated to both VIX options and SPX options simultaneously. One drawback of this model is that calibration may be slow because no closed form solution for European options exists. In this paper, we apply modified versions of the second order Monte Carlo scheme of Ninomiya and Victoir (2008) and compare these to the Euler-Maruyama scheme with full truncation of Lord et al. (2010), demonstrating on the one hand that fast calibration of the DMR model is practical, and on the other that suitably modified Ninomiya-Victoir schemes are applicable to the simulation of much more complicated time-homogeneous models than may have been thought previously.

1 Introduction

It is common knowledge that the Black-Scholes option pricing model is inconsistent with market pricing of options. Local volatility models, Lévy models, stochastic volatility models, stochastic volatility models with jumps and various variants and combinations of these have been proposed to fit market implied volatilities better and describe the dynamics of the resulting volatility surface. With the advent of trading in VIX options in 2007 however, marginal risk-neutral densities of forward volatilities of SPX became

*WIAS Berlin. christian.bayer@wias-berlin.de

†Department of Mathematics, Baruch College, CUNY.
jim.gatheral@baruch.cuny.edu

‡Department of Mathematical Sciences, University of Copenhagen.
m.karlsen@math.ku.dk

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. Notable amongst these are the market models of Bergomi (2005) and the variance curve factor models of Buehler (2006).

In Gatheral (2008), a specific three factor variance curve model was introduced with dynamics motivated by economic intuition for the empirical dynamics of the variance. This model was simultaneously calibrated to SPX and VIX option markets.

In this *double-mean-reverting* or *DMR* model, the dynamics are given by

$$dS_t = \sqrt{v_t} S_t dW_t^1, \quad (1.1a)$$

$$dv_t = \kappa_1 (v'_t - v_t) dt + \xi_1 v_t^{\alpha_1} dW_t^2, \quad (1.1b)$$

$$dv'_t = \kappa_2 (\theta - v'_t) dt + \xi_2 v_t^{\alpha_2} dW_t^3, \quad (1.1c)$$

where the Brownian motions W_i are all in general correlated with $\mathbb{E}[dW_t^i dW_t^j] = \rho_{ij} dt$.

Thus variance mean-reverts to a level that itself moves slowly over time with the state of the economy. Also, it is a stylized fact that the distribution of volatility (whether realized or implied) should be roughly lognormal (see Andersen et al. (2001) for example); when the model is calibrated to market option prices, we find that indeed $\alpha_1 \approx 1$ consistent with this stylized fact.

One drawback of this model is that 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 Gatheral (2008), the DMR model is calibrated using an Euler-Maruyama Monte Carlo scheme with the partial truncation step of Lord et al. (2010).

In this paper, we show how to apply the second order Monte Carlo scheme of Ninomiya and Victoir (2008) to the calibration of the DMR model, substantially improving calibration time. In passing, we show that a Ninomiya-Victoir second order Monte-Carlo scheme with fully closed-form steps can be achieved for models that are rather more complicated than those (such as the Heston model) to which the technique has been applied so far.

The plan of the paper is as follows. Section 2 describes how the model of Gatheral (2008) is calibrated. Section 3 explains the Monte Carlo scheme of Ninomiya and Victoir (2008), the drift trick of Bayer et al. (2013) and a subsequent extension by us which we apply to the DMR model. Section 4 presents practical examples of calibration to SPX and VIX options with

numerical results, and in Section 5 we perform a convergence analysis with reasonable parameters. In Section 6 we present some concluding remarks.

2 Estimating the constants of the DMR model

In Gatheral (2008), 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.

2.1 Estimation of κ_1 , κ_2 , θ and ρ_{23}

As of time t , the T -maturity forward variance is given by

$$\xi_t(T) = \mathbb{E}[v_T | \mathcal{F}_t]$$

and the T -maturity variance swap by

$$\mathbb{E} \left[\int_t^T v_s ds \middle| \mathcal{F}_t \right].$$

Variance swaps are traded in the market so forward variance is a traded asset. 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 Gatheral (2006) for example):

$$\mathbb{E} \left[\int_t^T v_s ds \middle| \mathcal{F}_t \right] = 2 \left\{ \int_{-\infty}^0 p(k) dk + \int_0^{\infty} c(k) dk \right\}, \quad (2.1)$$

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. Thus, given a database of historical market option prices, market variance swap prices may be estimated by interpolation, extrapolation and integration.

It is straightforward to verify that in the DMR model (1.1), forward variances are given by

$$\xi_t(T) = \theta + (v_t - \theta) e^{-\kappa_1 \tau} + (v'_t - \theta) \frac{\kappa_1}{\kappa_1 - \kappa_2} (e^{-\kappa_2 \tau} - e^{-\kappa_1 \tau}), \quad (2.2)$$

where $\tau = T - t$. Direct integration then gives us an expression for the spot

variance curve¹

$$\mathbb{E} \left[\int_t^T v_s ds \middle| \mathcal{F}_t \right] = \theta \tau + (v_t - \theta) \frac{1 - e^{-\kappa_1 \tau}}{\kappa_1} + (v'_t - \theta) \frac{\kappa_1}{\kappa_1 - \kappa_2} \left\{ \frac{1 - e^{-\kappa_2 \tau}}{\kappa_2} - \frac{1 - e^{-\kappa_1 \tau}}{\kappa_1} \right\}. \quad (2.3)$$

In the DMR model, θ , κ_1 and κ_2 are constants; v_t and v'_t are state variables. From (2.3), variance swaps depend linearly on the state variables. Thus, for every choice of θ , κ_1 and κ_2 , given option prices for various expiries, we can approximate the spot variance curve and infer v_t and v'_t by linear regression.

Performing daily regressions for the seven year period from January 2001 to April 2008, and optimizing over θ , κ_1 and κ_2 to minimize the mean squared error between the fitted curves and actual curves, optimal choices of the parameters θ , κ_1 and κ_2 and also daily time series of v_t and v'_t were obtained. The optimal choice of parameters was found to be

$$\theta = 0.078,$$

$$\kappa_1 = 5.5,$$

$$\kappa_2 = 0.10.$$

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

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

2.2 Estimation of the exponents α_1 and α_2

In order to obtain α_1 and α_2 we need information on how the volatility of volatility moves with the volatility itself. To obtain a proxy for the volatility of volatility Gatheral (2008) does the following.

Consider the SABR model for the forward with $\beta = 1$:

$$dF_t = \alpha_t F_t dW_t^1,$$

$$d\alpha_t = \nu \alpha_t dW_t^2,$$

¹(2.2) and (2.3) may be recognized as the Svensson yield curve model, re-used in our context.

where $dW_t^1 dW_t^2 = \rho dt$. An approximative Black-Scholes volatility for short maturities can be computed using the formula

$$\sigma_{BS}(k) = \alpha_0 f\left(\frac{k}{\alpha_0}\right) \quad (2.4)$$

where $k := \log(K/F_0)$ is the log-strike and

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

see Hagan et al. (2002). It is observed in Gatheral (2008) that the formula can fit observed volatilities very well, even for longer maturities. On a given day we have option quotes for a number of different maturities. We can fit the SABR model using the approximative formula (2.4) and obtain coefficients $\alpha_0^\tau, \nu^\tau, \rho^\tau$ for each maturity τ . Gatheral (2008) parametrizes the ν^τ coefficient for the different maturities with the function

$$\nu^\tau = \frac{\nu_{eff}}{\sqrt{\tau}}$$

which fits the term structure of the ν -parameter remarkably well. The number ν_{eff} is then used as a proxy for the volatility of volatility in a lognormal volatility model on that given day. The VIX index is used as a proxy for volatility.

Using the same dataset as in Section 2.1, Gatheral (2008) collects a time-series from January 2001 to April 2008 of ν_{eff} obtained by calibration to SPX options. He also collects VIX quotes. Doing a linear regression of $\log(\nu_{eff})$ onto $\log(\text{VIX})$ he obtains the equation

$$\log(\nu_{eff}) \approx -0.125 - 0.127 \log(\text{VIX}).$$

We can therefore write an SDE for the volatility:

$$d\alpha_t = c\alpha_t^{-0.127} \alpha_t dW_t = c\alpha_t^{0.873} dW_t^2.$$

In the DMR model we are looking for a coefficient on the variance $v_t = \alpha_t^2$. Using Ito's lemma we obtain

$$dv_t = O(dt) + 2cv_t^{0.9365} dW_t.$$

We will use the rounded coefficient $\alpha_1 = 0.94$, which obviously is close to one. There is insufficient market data to be able to say anything about the exponent α_2 so in Gatheral (2008), the choice $\alpha_2 = \alpha_1 = 0.94$ was made. As we will see in Section 3.2, various simplifications are possible if $\alpha_1 = \alpha_2 = 1$ (the so-called *double lognormal* model) so that case will also be considered in the following.

2.3 Daily calibration of remaining parameters

Although the volatility of volatility parameters ξ_1 and ξ_2 are in principle constants of the DMR model, Gatheral (2008) presents empirical evidence that calibrated parameters are not constant in the data. ξ_1 and ξ_2 are thus left free to be calibrated daily to VIX options data. The correlations ρ_{12} and ρ_{13} cannot be imputed from VIX option data; they are left free to improve the daily calibration of the DMR model to SPX data.

So on any given day, both the state variables v_t and v'_t , and the model parameters ξ_1 , ξ_2 , ρ_{12} and ρ_{13} are calibrated to VIX and SPX options data. v_t and v'_t are calibrated to variance swaps using linear regression and equation (2.3). In Gatheral (2008) calibration of ξ_1 , ξ_2 , ρ_{12} and ρ_{13} was performed using Monte-Carlo simulation. The chosen discretization was an Euler-Maruyama scheme with a partial truncation step, see Lord et al. (2010), which we can write recursively as

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

here 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 $\mathbb{E}[Z_k^i Z_k^j] = \rho_{ij}\Delta$. This is a general scheme and we do not need to know moments or asymptotic properties of the density in order to use it. Lord et al. (2010) finds the full truncation scheme superior when simulating the Heston model. When α_1 and α_2 are close to one however, our tests suggest that the partial truncation scheme is superior; this improvement becomes apparent only when time steps are large.

3 The Ninomiya-Victoir scheme and drift trick

In Ninomiya and Victoir (2008) a general second order weak discretization scheme for stochastic differential equations was introduced. Consider a multi-dimensional stochastic differential equation in Stratonovich form

$$d\mathbf{X}(t, \mathbf{x}) = V_0(\mathbf{X}(t, \mathbf{x}))dt + \sum_{i=1}^d V_i(\mathbf{X}(t, \mathbf{x})) \circ dB_t^i, \quad (3.1)$$

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

In this general setting, the Ninomiya-Victoir scheme based on a uniform grid with time steps Δ is recursively given by

$$\begin{aligned} \mathbf{X}^{(NV)}(0, \mathbf{x}) &= \mathbf{x}, \\ \mathbf{X}^{(NV)}((k+1)\Delta, \mathbf{x}) &= \begin{cases} e^{\frac{\Delta}{2}V_0} e^{Z_k^1 V_1} \dots e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}(k\Delta, \mathbf{x}), & \Lambda_k = -1, \\ e^{\frac{\Delta}{2}V_0} e^{Z_k^d V_d} \dots e^{Z_k^1 V_1} e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}(k\Delta, \mathbf{x}), & \Lambda_k = +1. \end{cases} \end{aligned} \quad (3.2)$$

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

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

i.e., the *flow* of the vector field V ,² and the probability space carries independent random-variables (Λ_k) , with values ± 1 at probability $1/2$, and independent $\mathcal{N}(0, \Delta)$ random variables (Z_k^j) . Note that $t = Z_k^j$ can take negative values, so one has to ensure that the ODE solutions used in an implementation of the NV scheme actually do make sense for positive as well as for negative t . One step in the NV scheme corresponds actually to a (non-discrete) cubature formula of order $m = 5$ in the sense of Lyons and Victoir (2004). When seen from this point of view, the reversal of the order of the flows depending on the coin-flip Λ_k serves to improve the approximation of the Lévy area and higher iterated integrals in the weak sense. On the other hand, one can also interpret the Ninomiya-Victoir scheme as the stochastic version of a classical operator splitting scheme, where the infinitesimal generator $V_0 + \frac{1}{2} \sum_{i=1}^d V_i^2$ of the diffusion is split into the first order differential operator V_0 and the second order differential operators $\frac{1}{2} V_1^2, \dots, \frac{1}{2} V_d^2$.³ The order-reversal is, in that context, a well known trick which improves the order of the method and goes back to Strang (1963).

The Ninomiya-Victoir scheme has attracted wide attention since its introduction; it is nowadays found in various sophisticated numerical packages such as Inria's software PREMIA for financial option computations. A variation of the scheme designed to deal with degeneracies arising in some affine situations is discussed in Alfonsi (2010).

² (3.2) is to be read from right to left, i.e., $e^{Z_k^d V_d} e^{\frac{\Delta}{2} V_0} \mathbf{X}^{(NV)}(k\Delta, \mathbf{x})$ means that solution $e^{\frac{\Delta}{2} V_0} \mathbf{X}^{(NV)}(k\Delta, \mathbf{x})$ of the ODE driven by V_0 is then used as initial value for the ODE driven by the vector field V_d , which is run until the (possibly negative) time Z_k^d .

³ Recall that a vector field $V : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is identified with the first order linear differential operator acting on smooth functions $f : \mathbb{R}^N \rightarrow \mathbb{R}$ by $Vf(x) \equiv \nabla f(x) \cdot V(x)$. By iteration, V^2 can then be interpreted as a linear second order differential operator.

3.1 Improving the efficiency of the Ninomiya-Victoir method

3.1.1 Changing the driving noise

In terms of numerical efficiency, cubature methods, and the Ninomiya-Victoir scheme in particular, heavily rely on the ability to solve, fast and accurately, ordinary differential equations. The general cubature methods involve *time-inhomogeneous* ODEs with a rather complicated structure, involving all vector-fields at all times. Thus, there is usually no alternative to solving them numerically, often with Runge-Kutta methods. (A detailed discussion on how Runge-Kutta methods are applied in this context is found in Ninomiya and Ninomiya (2009).)

Using the canonical splitting induced by the model formulation, the Ninomiya-Victoir scheme only involves the composition of solution flows to time-homogeneous ODEs. In particular, there will be “lucky” cases of models where all (or at least most) ODE flows can 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, and the Ninomiya-Victoir method can be expected to perform especially well in such cases. As was observed by Ninomiya and Victoir (2008), the Heston model is such a lucky case. However, one soon encounters models (e.g., the popular SABR model) in which some of the vector-fields do not allow for flows in closed form. In Bayer et al. (2013), it was found that the class of favorable models can be significantly enlarged by working with an almost trivial modification of the NV scheme. This modification is based on the equivalence of (3.1) with

$$\begin{aligned} d\mathbf{X}(t, \mathbf{x}) &= \left(V_0(\mathbf{X}(t, \mathbf{x})) - \sum_{j=1}^d \gamma_j V_j(\mathbf{X}(t, \mathbf{x})) \right) dt + \\ &\quad + \sum_{j=1}^d V_j(\mathbf{X}(t, \mathbf{x})) \circ d \left(B_t^j + \gamma_j t \right) \\ &\equiv V_0^{(\gamma)}(\mathbf{X}(t, \mathbf{x})) dt + \sum_{j=1}^d V_j(\mathbf{X}(t, \mathbf{x})) \circ d \left(B_t^j + \gamma_j t \right) \end{aligned}$$

whatever the choice of drift parameters $\gamma_1, \dots, \gamma_d$. Assume that all diffusion vector-fields (V_1, \dots, V_d) allow for flows in closed form, whereas e^{tV_0} is not available in closed form.⁴ The point is that, in a variety of concrete examples,

⁴As models are usually devised in the Ito framework and V_0 is obtained from the Ito drift by the Stratonovich correction, this situation is quite common in finance.

one can pick drift parameters $\gamma_1, \dots, \gamma_d$ in a way that $e^{tV_0^{(\gamma)}}$ can be solved in closed form after all.

Therefore, we propose the following variant of the Ninomiya-Victoir method (which, following Bayer et al. (2013), shall be referred to as the ‘‘Ninomiya-Victoir scheme with drift (trick)’’):

$$\begin{aligned} \mathbf{X}^{(NVd)}(0, \mathbf{x}) &= \mathbf{x}, \\ \mathbf{X}^{(NVd)}((k+1)\Delta, \mathbf{x}) &= \\ &\begin{cases} e^{\frac{\Delta}{2}V_0^{(\gamma)}} e^{Z_k^1 V_1} \dots e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_0^{(\gamma)}} \mathbf{X}^{(NVd)}(k\Delta, \mathbf{x}), & \Lambda_k = -1, \\ e^{\frac{\Delta}{2}V_0^{(\gamma)}} e^{Z_k^d V_d} \dots e^{Z_k^1 V_1} e^{\frac{\Delta}{2}V_0^{(\gamma)}} \mathbf{X}^{(NVd)}(k\Delta, \mathbf{x}), & \Lambda_k = +1, \end{cases} \end{aligned} \quad (3.3)$$

where $Z_k^i \sim \mathcal{N}(\Delta\gamma_i, \Delta)$ independent of each other.

Note that (3.3) corresponds to the splitting of the differential operator according to

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

3.1.2 Incorporating ODE splitting

The strategy of Section 3.1.1, namely to replace the standard Ninomiya-Victoir splitting (3.2) by a different one customized to the specific problem at hand, can be generalized to accommodate for an even wider class of problems. In particular, one can directly incorporate any *splitting scheme* (in the ODE sense) for any of the ODEs involved in (3.2) into the Ninomiya-Victoir scheme. Let us again assume that (only) the Stratonovich vector field V_0 is too complicated to allow for closed form solutions of the corresponding ODEs. The structure of the Stratonovich drift vector field

$$V_0(\mathbf{x}) = V(\mathbf{x}) - \frac{1}{2} \sum_{i=1}^d DV_i(\mathbf{x}) \cdot V_i(\mathbf{x}),$$

where $V : \mathbb{R}^N \rightarrow \mathbb{R}^N$ denotes the drift vector field of the SDE in the Ito formulation and DV_i denotes the Jacobian of the vector field V_i , $i = 1, \dots, d$, motivates to apply a classical ODE splitting scheme in order to solve the ODE $\dot{\mathbf{y}} = V_0(\mathbf{y})$, i.e., we 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

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

a method sometimes known as *symplectic Euler scheme*, see Hairer et al. (2006). We can incorporate the symplectic Euler method in the Ninomiya-Victoir scheme as follows: starting with $\mathbf{X}^{(NVs)}(0, \mathbf{x}) = \mathbf{x}$, we iterate according to

$$\begin{aligned} \mathbf{X}^{(NVs)}((k+1)\Delta, \mathbf{x}) = & \\ \begin{cases} e^{\frac{\Delta}{2}V_{0,1}}e^{\frac{\Delta}{2}V_{0,2}}e^{Z_k^1V_1} \dots e^{Z_k^dV_d}e^{\frac{\Delta}{2}V_{0,2}}e^{\frac{\Delta}{2}V_{0,1}}\mathbf{X}^{(NVs)}(k\Delta, \mathbf{x}), & \Lambda_k = -1, \\ e^{\frac{\Delta}{2}V_{0,1}}e^{\frac{\Delta}{2}V_{0,2}}e^{Z_k^dV_d} \dots e^{Z_k^1V_1}e^{\frac{\Delta}{2}V_{0,2}}e^{\frac{\Delta}{2}V_{0,1}}\mathbf{X}^{(NVs)}(k\Delta, \mathbf{x}), & \Lambda_k = +1. \end{cases} \end{aligned} \quad (3.4)$$

Even though the symplectic Euler scheme only has local order two, the Strang trick of repeating the symplectic Euler scheme once while inverting the order of the vector fields, again produces a scheme with local order three and, hence, global order two. Indeed, note that the *Verlet scheme*

$$e^{\Delta V_0} \mathbf{x} = e^{\frac{\Delta}{2}V_{0,1}}e^{\Delta V_{0,2}}e^{\frac{\Delta}{2}V_{0,1}}\mathbf{x} + \mathcal{O}(\Delta^3)$$

obtained by omitting the diffusion part in (3.4) has (global) order two. Both the Verlet and the symplectic Euler scheme are examples of *geometric integrators* for ODEs, and we again refer to Hairer et al. (2006) for much more information.

3.1.3 Analysis of the modified Ninomiya Victoir scheme of (3.4)

We give a sketch of the proof that the modified NV algorithm (3.4) has second order convergence in the weak sense. In particular, in the following we assume sufficient regularity for all involved functions and vector fields. Let \mathbf{X}_Δ denote the true solution of the SDE at time Δ and let $\bar{\mathbf{X}}_\Delta \equiv \mathbf{X}^{(NVs)}(\Delta, \mathbf{x})$ denote the output of the modified NV algorithm after one timestep of size Δ , both started at \mathbf{x} at time 0. By the Markov property, it suffices to show that the weak local error is of third order (see, for instance, Talay and Tubaro (1990)), i.e., that

$$\mathbb{E}[f(\mathbf{X}_\Delta)] - \mathbb{E}[f(\bar{\mathbf{X}}_\Delta)] = \mathcal{O}(\Delta^3) \quad (3.5)$$

for sufficiently smooth test functions f . Indeed, denote $u(t, \mathbf{y}) = \mathbb{E}[g(\mathbf{X}_T) | \mathbf{X}_t = \mathbf{y}]$ and assume that we want to approximate $\mathbb{E}[g(\mathbf{X}_T)] = u(0, \mathbf{x})$ by $\mathbb{E}[g(\bar{\mathbf{X}}_{n\Delta})] = \mathbb{E}[u(n\Delta, \bar{\mathbf{X}}_{n\Delta})]$ with $\Delta = T/n$. Then, by a telescoping sum, we may decompose the global error $\mathbb{E}[g(\bar{\mathbf{X}}_{n\Delta})] - u(0, \mathbf{x})$ as the sum of the local errors

$$\mathbb{E}[g(\bar{\mathbf{X}}_{n\Delta})] - u(0, \mathbf{x}) = \sum_{k=0}^{n-1} \mathbb{E}[u((k+1)\Delta, \bar{\mathbf{X}}_{(k+1)\Delta}) - u(k\Delta, \bar{\mathbf{X}}_{k\Delta})].$$

Assuming (3.5), we have

$$\mathbb{E} [u((k+1)\Delta, \bar{\mathbf{X}}_{(k+1)\Delta}) - u(k\Delta, \bar{\mathbf{X}}_{k\Delta})] = \mathcal{O}(\Delta^3),$$

by first conditioning on $\bar{\mathbf{X}}_{k\Delta}$. However, we sum $n = T/\Delta$ of these terms, so that the global error is $\mathcal{O}(\Delta^2)$.

For the proof of (3.5), note that the ‘‘multiplication’’ of vector fields in the sense of iterative applications of vector fields as differential operators is certainly non-commutative. Taylor expansion applied to (3.4) implies that

$$\begin{aligned} \mathbb{E}[f(\bar{\mathbf{X}}_\Delta)] &= \frac{1}{2} \left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2 \right) \left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2 \right) \\ &\quad \left(1 + \frac{1}{2}\Delta V_1^2 + \frac{1}{8}\Delta^2 V_1^4 \right) \cdots \left(1 + \frac{1}{2}\Delta V_d^2 + \frac{1}{8}\Delta^2 V_d^4 \right) \\ &\quad \left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2 \right) \left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2 \right) f(\mathbf{x}) + \\ &\quad + \frac{1}{2} \left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2 \right) \left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2 \right) \\ &\quad \left(1 + \frac{1}{2}\Delta V_d^2 + \frac{1}{8}\Delta^2 V_d^4 \right) \cdots \left(1 + \frac{1}{2}\Delta V_1^2 + \frac{1}{8}\Delta^2 V_1^4 \right) \\ &\quad \left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2 \right) \left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2 \right) f(\mathbf{x}) + \mathcal{O}(\Delta^3) \\ &= f(\mathbf{x}) + \Delta \left(V_0 + \frac{1}{2} \sum_{i=1}^d V_i^2 \right) f(\mathbf{x}) + \\ &\quad + \frac{1}{2} \Delta^2 \left(V_0 + \frac{1}{2} \sum_{i=1}^d V_i^2 \right)^2 f(\mathbf{x}) + \mathcal{O}(\Delta^3) \\ &= \mathbb{E}[f(\mathbf{X}_\Delta)] + \mathcal{O}(\Delta^3), \end{aligned}$$

where the last equality follows since $V_0 + \frac{1}{2} \sum_{i=1}^d V_i^2$ is the infinitesimal generator of the diffusion $\mathbf{X}(t, \mathbf{x})$.

3.2 The Ninomiya-Victoir scheme for the DMR model

3.2.1 The Stratonovich formulation of the DMR model

Consider again the DMR model (1.1) re-expressed in terms of independent Brownian motions B^i :

$$\begin{aligned} dS_t &= \sqrt{v_t} S_t dB_t^1, \\ dv_t &= \kappa_1 (v'_t - v_t) dt + \xi_1 v_t^{\alpha_1} \left(\tilde{\rho}_{1,2} dB_t^1 + \sqrt{1 - \tilde{\rho}_{1,2}^2} dB_t^2 \right), \\ dv'_t &= \kappa_2 (\theta - v'_t) dt + \xi_2 v_t'^{\alpha_2} \left(\tilde{\rho}_{1,3} dB_t^1 + \tilde{\rho}_{2,3} dB_t^2 + \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2} dB_t^3 \right), \end{aligned} \quad (3.6)$$

where $\tilde{\rho}_{12} = \rho_{12}$, $\tilde{\rho}_{13} = \rho_{13}$ and $\tilde{\rho}_{23} = \frac{\rho_{23} - \rho_{12}\rho_{13}}{\sqrt{1 - \rho_{12}^2}}$. In order to apply the simulation method of Ninomiya and Victoir (2008) we need to re-express the Itô SDEs (3.6) in Stratonovich form (see for example Definition 3.13 of Karatzas and Shreve (1988)).

We first compute the quadratic covariation terms as follows:

$$\begin{aligned} d[\sqrt{v_t} S, B^1]_t &= \left\{ \frac{1}{2} \tilde{\rho}_{1,2} v_t^{\alpha_1 - \frac{1}{2}} + v_t \right\} S_t dt \\ d\left[\xi_1 v_t^{\alpha_1}, \left(\tilde{\rho}_{1,2} B^1 + \sqrt{1 - \tilde{\rho}_{1,2}^2} B^2 \right) \right]_t &= \xi_1^2 \alpha_1 v_t^{2\alpha_1 - 1} dt \\ d\left[\xi_2 v_t'^{\alpha_2}, \left(\tilde{\rho}_{1,3} B^1 + \tilde{\rho}_{2,3} B^2 + \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2} B^3 \right) \right]_t &= \xi_2^2 \alpha_2 v_t'^{2\alpha_2 - 1} dt. \end{aligned}$$

We then obtain the Stratonovich form of (3.6):

$$\mathbf{X}(t, \mathbf{x}) = \mathbf{x} + \int_0^t V_0(\mathbf{X}(s, \mathbf{x})) ds + \sum_{j=1}^3 \int_0^t V_j(\mathbf{X}(s, \mathbf{x})) \circ dB_s^j \quad (3.7)$$

where the state vector $\mathbf{X}(t, \mathbf{x}) = (S_t, v_t, v'_t)^T$, the initial condition $\mathbf{x} = (S_0, v_0, v'_0)^T$, and the driving vector fields are given by

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

and

$$\begin{aligned} V_1(\mathbf{x}) &= (\sqrt{x_2} x_1 \quad \tilde{\rho}_{1,2} \xi_1 x_2^{\alpha_1} \quad \tilde{\rho}_{1,3} \xi_2 x_3^{\alpha_2})^T, \\ V_2(\mathbf{x}) &= \left(0 \quad \sqrt{1 - \tilde{\rho}_{1,2}^2} \xi_1 x_2^{\alpha_1} \quad \tilde{\rho}_{2,3} \xi_2 x_3^{\alpha_2} \right)^T, \\ V_3(\mathbf{x}) &= \left(0 \quad 0 \quad \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2} \xi_2 x_3^{\alpha_2} \right)^T. \end{aligned}$$

In order to proceed with the Ninomiya-Victoir splitting, we thus need to solve the ordinary differential equations

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

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

3.2.2 The flow of the diffusion vector fields

Following Bayer et al. (2013), it is straightforward to verify that the solution to the (one-dimensional) ODE

$$\frac{d}{dt} x(t) = h(t)^\alpha x(t)^\beta \tag{3.8}$$

is given by

$$x(t) = \begin{cases} [(1 - \beta)H(t) + x(0)^{1-\beta}]_+^{\frac{1}{1-\beta}}, & 0 < \beta < 1, \\ x(0) e^{H(t)}, & \beta = 1, \end{cases} \tag{3.9}$$

where

$$H(t) = \int_0^t h(s)^\alpha ds.$$

We can thus solve the ODEs

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

for $i \in \{1, 2, 3\}$ in closed form. The NV algorithm requires solutions of the ODEs driven by the diffusion vector fields for negative times t . As $\frac{d}{dt} x(-t) = -\dot{x}(-t)$, this essentially means that the sign of the coefficient changes. In any case, for a fixed time interval $I = [0, T]$ (or $I = [-T, 0]$ in the negative time case), (3.9) is the unique solution to (3.8), provided that $x(0) \neq 0$. This follows by standard arguments when $x > 0$ on I . On the other hand, note that in our case $h(t)$ cannot change its sign on I for the

ODEs under considerations here. Thus, H is always a monotonous function on I . So, for $x(0) > 0$, we can only get $x(t) = 0$ for some $t \in I$, if H is negative and decreasing (increasing for the negative time case). But then x must stay at 0 for the remaining time to T (or $-T$, respectively) – as we do not allow for complex-valued solutions. If, however, $x(0) = 0$, then there might, indeed, be several real-valued solutions for (3.8). For instance in the positive time case, when $h > 0$, both (3.9) and $x \equiv 0$ are solutions. However, when paths of the underlying SDE come too close to 0 too often, then the NV scheme should not be expected to perform any better than the EM scheme, as has been found in several numerical studies, e.g., by Lord et al. (2010). Indeed, the whole theory of NV breaks actually down in that case.

3.2.3 The flow of the Stratonovich drift vector field

Solving the ODE for $i = 0$ is a little trickier with no obvious closed-form solution for general values of the exponents α_i . Whereas a numerical solution would be possible by for example applying a Runge-Kutta method, by further splitting the operator, we may obtain a reasonably simple closed-form simulation step. That is, we write

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

with

$$V_{0,1}(x) = \begin{pmatrix} -\frac{1}{2} x_2 x_1 \\ -\kappa_1 (x_2 - x_3) \\ -\kappa_2 (x_3 - \theta) \end{pmatrix}, \quad V_{0,2}(x) = \begin{pmatrix} -\frac{1}{4} \xi_1 \tilde{\rho}_{1,2} x_2^{\alpha_1 - \frac{1}{2}} x_1 \\ -\frac{1}{2} \xi_1^2 \alpha_1 x_2^{2\alpha_1 - 1} \\ -\frac{1}{2} \xi_2^2 \alpha_2 x_3^{2\alpha_2 - 1} \end{pmatrix}.$$

and solve the equations

$$\frac{d}{dt} x(t) = V_{0,j}(x(t)) \text{ with } j = 1, 2.$$

Solution for $j = 1$

The equation in the third row which reads

$$\frac{d}{dt} x_3(t) = -\kappa_2 (x_3 - \theta)$$

has the solution

$$x_3(t) = \theta + e^{-\kappa_2 t} (x_3(0) - \theta).$$

The second ODE has the solution

$$x_2(t) = e^{-\kappa_1 t} x_2(0) + \kappa_1 \int_0^t e^{-\kappa_1 (t-s)} x_3(s) ds$$

which is just the forward variance curve $\xi_t(T)$. The first ODE reads

$$\frac{d}{dt} x_1(t) = -\frac{1}{2} x_2(t) x_1(t)$$

with the solution

$$x_1(t) = x_1(0) \exp \left\{ -\frac{1}{2} H(t) \right\}$$

with $H(t) = \int_0^t x_2(s) ds$ which we recognize as the spot variance swap curve.

Solution for $j = 2$

If $\alpha_1 \neq 1$ and $\alpha_2 \neq 1$, the second and third ODEs have solutions

$$\begin{aligned} x_2(t) &= \left[x_2(0)^{2(1-\alpha_1)} - \alpha_1 (1 - \alpha_1) \xi_1^2 t \right]_+^{\frac{1}{2(1-\alpha_1)}}, \\ x_3(t) &= \left[x_3(0)^{2(1-\alpha_2)} - \alpha_2 (1 - \alpha_2) \xi_2^2 t \right]_+^{\frac{1}{2(1-\alpha_2)}}. \end{aligned}$$

The first ODE reads

$$\frac{d}{dt} x_1(t) = -\frac{1}{4} \xi_1 \tilde{\rho}_{1,2} x_2^{\alpha_1 - \frac{1}{2}} x_1$$

with the solution

$$x_1(t) = x_1(0) \exp \left\{ -\frac{1}{4} \xi_1 \tilde{\rho}_{1,2} H(t) \right\}$$

where $H(t) = \int_0^t x_2(s)^{\alpha_1 - \frac{1}{2}} ds$. $H(t)$ can also then be computed explicitly as

$$H(t) = \frac{2}{\alpha_1 (3/2 - \alpha_1) \xi_1^2} \left\{ x_2(0)^{3/2 - \alpha_1} - x_2(t)^{3/2 - \alpha_1} \right\}.$$

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

The special case $\alpha_1 = 1$, $\alpha_2 = 1$ gives the double lognormal model of Gatheral (2008), a model which both fits the empirical SPX and VIX surfaces well and displays remarkable parameter stability. In this case, we may obtain even simpler closed-form solutions by applying the drift trick of Bayer et al. (2013)

explained in Section 3. This trick involves simplifying V_0 by subtracting components spanned by the the vector fields V_1, V_2, V_3 . This is achieved by introducing drift in the Brownian motions. Specifically, with

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

and choosing

$$\begin{aligned}\gamma_1 &= -\xi_1 \tilde{\rho}_{1,2}, \\ \gamma_2 &= -\frac{\kappa_1 + \frac{1}{2}\xi_1^2 + \gamma_1 \tilde{\rho}_{1,2} \xi_1}{\xi_1 \sqrt{1 - \tilde{\rho}_{1,2}^2}}, \\ \gamma_3 &= -\frac{\kappa_2 + \frac{1}{2}\xi_2^2 - \tilde{\rho}_{1,3} \xi_2 \gamma_1 - \tilde{\rho}_{2,3} \xi_2 \gamma_2}{\xi_2 \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2}},\end{aligned}$$

we have the much simpler expression

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

The third and second differential equations respectively have solutions

$$\begin{aligned}x_3(t) &= x_3(0) + \kappa_2 \theta t, \\ x_2(t) &= x_2(0) + \kappa_1 \left(x_3(0) t + \frac{1}{2} \kappa_2 \theta t^2 \right).\end{aligned}$$

The first ODE reads

$$\frac{d}{dt} x_1(t) = -\frac{1}{2} x_2(t) x_1(t)$$

with the solution

$$x_1(t) = x_1(0) \exp \left\{ -\frac{1}{2} H(t) \right\}$$

where

$$\begin{aligned}H(t) &= \int_0^t x_2(s) ds \\ &= x_2(0) t + \kappa_1 \left(\frac{1}{2} x_3(0) t + \frac{1}{6} \kappa_2 \theta t^3 \right).\end{aligned}$$

3.3 Summary of the modified Ninomiya-Victoir procedure

Adopting the notation of Bayer et al. (2013), if Λ_k is a Bernoulli random variable, the k th NV time step of length Δ in the modified NV simulation of Section 3.2.3 is of the form

$$\begin{aligned} & \mathbf{X}((k+1)\Delta, \mathbf{x}) \\ = & \begin{cases} e^{\frac{1}{2}\Delta V_{0,1}} e^{\frac{1}{2}\Delta V_{0,2}} e^{Z_k^1 V_1} e^{Z_k^2 V_2} e^{Z_k^3 V_3} e^{\frac{1}{2}\Delta V_{0,2}} e^{\frac{1}{2}\Delta V_{0,1}} \mathbf{X}(k\Delta, \mathbf{x}) & \text{if } \Lambda_k = -1, \\ e^{\frac{1}{2}\Delta V_{0,1}} e^{\frac{1}{2}\Delta V_{0,2}} e^{Z_k^3 V_3} e^{Z_k^2 V_2} e^{Z_k^1 V_1} e^{\frac{1}{2}\Delta V_{0,2}} e^{\frac{1}{2}\Delta V_{0,1}} \mathbf{X}(k\Delta, \mathbf{x}) & \text{if } \Lambda_k = +1 \end{cases} \end{aligned}$$

where the $Z_k^i \sim N(0, \Delta)$ are independent of each other.

Similarly, in the NV procedure with drift trick of Section 3.2.4, the k th NV time step is of the simpler form

$$\begin{aligned} & \mathbf{X}((k+1)\Delta, \mathbf{x}) \\ = & \begin{cases} e^{\frac{1}{2}\Delta V_0^\gamma} e^{Z_k^1 V_1} e^{Z_k^2 V_2} e^{Z_k^3 V_3} e^{\frac{1}{2}\Delta V_0^\gamma} \mathbf{X}(k\Delta, \mathbf{x}) & \text{if } \Lambda_k = -1 \\ e^{\frac{1}{2}\Delta V_0^\gamma} e^{Z_k^3 V_3} e^{Z_k^2 V_2} e^{Z_k^1 V_1} e^{\frac{1}{2}\Delta V_0^\gamma} \mathbf{X}(k\Delta, \mathbf{x}) & \text{if } \Lambda_k = +1 \end{cases} \end{aligned}$$

where the $Z_k^i \sim N(\gamma_i \Delta, \Delta)$ (note the nonzero drift) are independent of each other.

4 Calibrating the model

4.1 Calibrating daily parameters

As mentioned in Section 2.3 we would like to infer $v_t, v'_t, \xi_1, \xi_2, \rho_{12}$ and ρ_{13} daily. The calibration of these parameters is divided into several steps.

4.1.1 v_0 and v'_0

We saw in Section 2.1 that the prices of variance swaps may be estimated from the market prices of SPX options using equation (2.1). This computation requires a continuous set of option prices which we obtain by fitting the SVI parametrization (see for example Gatheral and Jacquier (2012)) to the volatility smile for each expiry. Then from (2.3), given κ_1, κ_2 and θ , v_0 and v'_0 may be obtained by linear regression.

4.1.2 ξ_1 and ξ_2

The volatility parameters ξ_1 and ξ_2 of the variance processes are obtained by calibrating the model to the market prices of VIX options.

We proxy the underlying of a VIX option by the expected forward variance in our model. The payoff of a call option on the VIX index with strike K expiring at time T may therefore be written as

$$\left(\sqrt{\mathbb{E} \left[\int_T^{T+\Delta} v_s ds \mid \mathcal{F}_T \right]} - K \right)^+$$

where Δ is the length (approximately one month) of the VIX index. For each Monte Carlo path we have a value for v_T and v'_T , so the expected forward variance $\mathbb{E} \left[\int_T^{T+\Delta} v_s ds \mid \mathcal{F}_T \right]$ may be computed using (2.3). Averaging over all paths gives the model price of the VIX option.

Our chosen objective function is the sum of squared difference between market VIX option prices and the model VIX option prices, both expressed in terms of Black-Scholes implied volatility. Errors are weighted by the reciprocal of the bid-ask spread:

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

The minimization is performed with a Levenberg-Marquardt algorithm, setting starting values to $\xi_1 = 2.5$ and $\xi_2 = 0.4$, values typical of those that we observe.

4.1.3 ρ_{12} and ρ_{13}

We are then left with the two parameters ρ_{12} and ρ_{13} to calibrate. These are used to fit the SPX volatility surface. Our chosen objective function is again the sum of squared differences between market SPX option prices and model SPX option prices, all in implied volatility terms and weighted by the reciprocal of the bid-ask spread.

The objective function can though have multiple local minima, which lead to poor performance of the Levenberg-Marquardt algorithm. This is especially true when using the EM algorithm in combination with pseudo random numbers. To improve performance we need to find a good starting point before applying the solver. We achieve this by evaluating the function at a number of points and starting the Levenberg-Marquardt algorithm at the best point.

When we fit the Heston model to SPX option data, the imputed stock-volatility correlation parameter is typically around -0.7 . It seems reasonable

then that the two correlation parameters ρ_{12} and ρ_{13} should be in the same ballpark. We therefore search for ρ_{12} and ρ_{13} in the region $[-1, -0.5] \times [-1, -0.5]$, restricted by condition that $\tilde{\rho}_{13}^2 + \tilde{\rho}_{23}^2 \leq 1$. A good starting point can be found by evaluating the objective function at 30 Sobol points in this region.

4.2 Calibration examples with tests of Monte Carlo schemes

We now consider two calibration examples: One with data from April 3, 2007, and the other with data from September 15, 2011. In both of these examples, we will compare the calibration performance of the modified Ninomiya-Victoir scheme described in Section 3 with that of the Euler-Maruyama scheme with partial truncation step described in Section 2.3. We thereby test both the model and the calibration routines in pre- and post-crisis environments.

Our testing strategy is as follows: With respectively $2^9 = 512$, $2^{11} = 2048$ and $2^{13} = 8192$ paths we calibrate the model using 6, 10, 20, \dots , 100, 200, 300, 400, 500, 1000, 2000 time steps, thereby obtaining calibrated values for ξ_1 and ξ_2 .

For each such calibration we obtain optimal volatility parameters. For example $\xi_1^{opt,30,11}$ and $\xi_1^{opt,30,11}$, are optimal parameters for a calibration with 30 time steps and 2^{11} paths. Using the optimal parameters we then perform another Ninomiya-Victoir Monte Carlo simulation with $2^{16} = 65,536$ paths and 500 timesteps. This latter simulation we use to measure how well the parameters obtained by the calibration fit the market. This is done by calculating the mean squared error objective function (RMSE) between the model prices obtained using $2^{16} = 65,536$ paths, 500 timesteps and the market prices. We can then assess what the minimum required number of paths and time steps is to obtain a calibration accurate enough for practical applications.

We also compare the performance of a classic Monte Carlo (MC) scheme using pseudo random numbers with that of a Quasi Monte Carlo (QMC) scheme using Sobol quasi random numbers. If the dimension of the integration problem is small, Quasi Monte Carlo should theoretically result in a lower integration error compared to Monte Carlo. Increasing the number of dimensions however decreases the efficiency of the Quasi Monte Carlo method and at some point the Monte Carlo method beats it with a lower integration error. In our implementation, the dimension of the integration problem is the number of time steps times the number of random variables required per timestep (which is three for EM and four for NV due to the

coin flip).

There exists a number of heuristic ways to deal with the curse of dimensionality of QMC, see for example da Silva and Barbe (2005). These methods will though not be tried out.

All computations were run on an ASUS desktop with an Intel Core i3 cpu at 2.40 GHz CPU and 4GB memory. The simulations were done in Java using the SSJ package, see “<http://www.iro.umontreal.ca/~simardr/ssj/indexe.html>”. For the Sobol sequences we used the built in direction numbers up to 360 dimensions. Sequences with more dimension were created using direction numbers from the webpage “<http://web.maths.unsw.edu.au/~fkuo/sobol/new-joe-kuo-6.21201>”, these have been obtained using the search criteria $D^{(6)}$ see Joe and Kou (2008). Pseudo random numbers were generated using the Mersenne twister, MT19937. Optimization were done using the Levenberg-Marquardt algorithm present in the SSJ package, this is a Java translation of the MINPACK routine, see More et al. (1980).

4.3 April 3, 2007

4.3.1 The data

The SPX option dataset contains prices for 421 options, 388 of them include both bid and ask prices, we only use these options in our calibration. There are 14 different option maturities in the dataset ranging from 0.005 to 2.71 years, the forward for the first maturity is 1438.62 and for the last maturity 1556.75. The strikes for the different options lie in the interval 600 to 2000. The VIX option dataset contains prices for 108 options, 96 of them include both bid and ask prices, again we only use options with both bid and ask prices. The dataset contains 7 different maturities ranging from 0.04 years to 1.13 years. The forward for the first maturity is 13.97 and 15.29 for the last maturity. Strikes lie in the interval 10 to 30.

4.3.2 Calibration of v_0 and v'_0

As explained in Section 4.1.1, we use linear regression to calibrate model variance swaps to market variance swaps (proxied by the SVI log-strip) giving us the parameters

$$\begin{aligned}v_0 &= 0.0153, \\v'_0 &= 0.0224.\end{aligned}$$

The resulting fit is graphed in Figure 1.

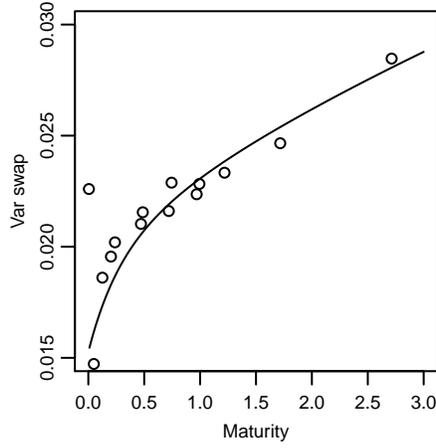


Figure 1: SPX market variance swaps as points together with the calibrated model variance swap curve (solid line). Data from April 3, 2007.

4.3.3 Calibration of ξ_1 and ξ_2 to VIX options

In Figure 2 we have graphed the RMSE from the different calibrations as a function of number of timesteps used.

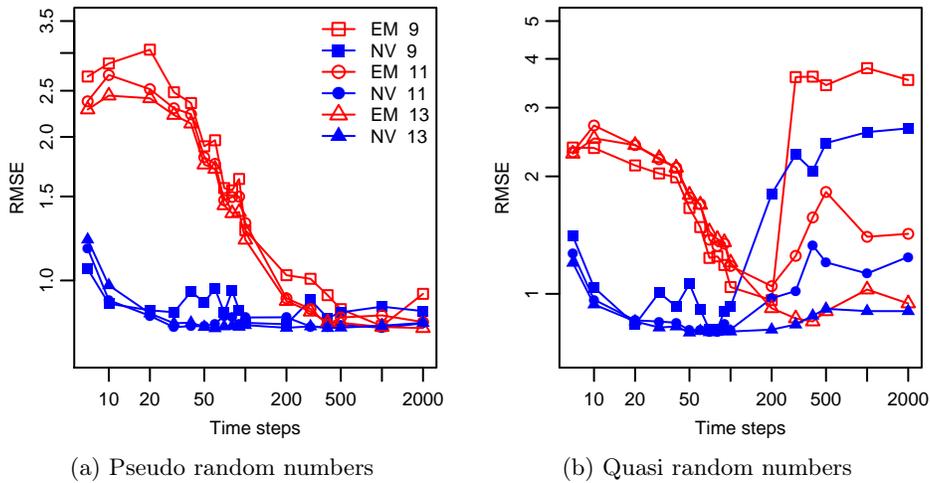


Figure 2: RMSE from NV and EM calibration of the DMR model to VIX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (a) specifies the method and the \log_2 number of paths.

The figure clearly indicate that a lower bound for the calibration RMSE exists around 0.8. Even with 2^{13} paths we do not get below this barrier. We see a clear advantage of using the NV scheme compared to the EM scheme. Using the NV scheme a calibration can be done using 20 – 30 timesteps and 2^{11} or 2^{13} paths. In comparison we need 400 – 500 timesteps if we use the EM scheme. Therefore we can reduce the number of timesteps by a factor of 15 or so.

The integration error seems to be negligible at 2^{11} paths. A good compromise between calibration quality and computational cost therefore seems to be an NV scheme using 2^{11} paths and 30 timesteps.

In Figure 3 we graph market VIX Black-Scholes implied volatility smiles together with model smiles. The model parameters ξ_1 and ξ_2 were calibrated to the market using MC with 30 NV time steps and 2^{11} paths. Total calibration time was 1.47 seconds. The resulting calibrated ξ parameters are:

$$\begin{aligned}\xi_1 &= 2.873, \\ \xi_2 &= 0.302.\end{aligned}$$

Model option prices were then computed using 100 NV QMC time steps and 2^{16} paths. By inspection of Figure 3, the quality of the calibration is quite acceptable, though VIX option bid-ask spreads are admittedly wide.

4.3.4 Calibration of ρ_{12} and ρ_{13} to SPX options

This test consists in fixing all parameters including the values of ξ_1 and ξ_2 found by the MC-NV calibration of section 4.3.3.

Let us start by doing a normal Levenberg-Marquadt calibration started in $(-0.7, -0.7)$, without an initial search for an optimal point. Figure 4 shows the RMSE results

The NV scheme performs acceptably, especially using 2^{13} paths, but the 2^9 paths calibration seems to be unacceptably off. The EM scheme shows strange behavior: note that for the first number of time steps the 2^9 paths calibration does best of all the methods, while the EM scheme with more paths shows jumpy behavior, sometimes it finds a good minimum other times it does not. In order to improve the method we use the simple search algorithm described in Section 4.1.3. In Figure 5 we have applied this before the Levenberg-Marquadt optimizer. After the search algorithm has been applied the EM scheme performs just as well if not better than the NV scheme. The NV scheme with 2^{13} paths and 30 time steps could in principle save us from using the search algorithm but since the NV scheme

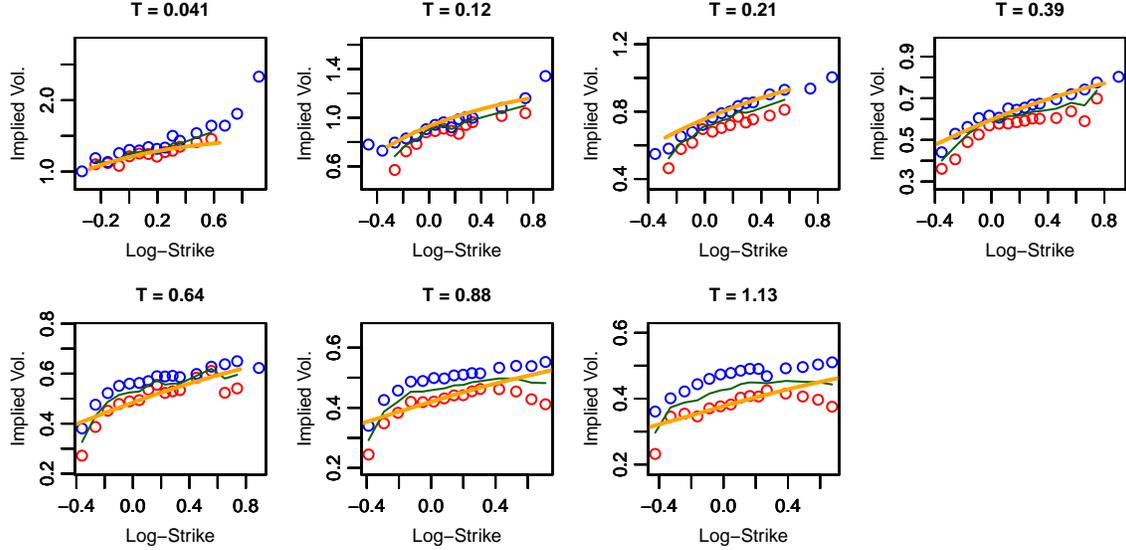


Figure 3: Implied Black volatilities for VIX options on April 3, 2007 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from a QMC-NV scheme using 100 time steps and 2^{16} paths (yellow line). The model parameters ξ_1 and ξ_2 are obtained by a calibration using a MC-NV scheme with 30 time steps and 2^{11} Monte Carlo paths.

is much more involved than the EM scheme we will not obtain a speedup. The RMSEs reported in Figure 5 clearly show that 2^{11} QMC paths are sufficient in order to obtain a good calibration. The spikes in figure 5(a) seems strange, but they only exists for a low number of paths and for pseudo random number. We conclude that it is best to use quasi random numbers when calibrating to SPX options.

In Figure 6, we graph market SPX Black-Scholes implied volatility smiles together with model smiles. The correlation parameters ρ_{12} and ρ_{13} were calibrated to the market using QMC with 30 EM time steps and 2^{11} paths. Total calibration time was 2.94 seconds in this case. The resulting calibrated parameters are:

$$\begin{aligned}\rho_{12} &= -0.992, \\ \rho_{13} &= -0.615\end{aligned}$$

Model SPX option prices were again computed using a QMC-NV scheme with 100 time steps and 2^{16} paths.

Inspecting Figure 6, we note that the DMR model fits longer expiration

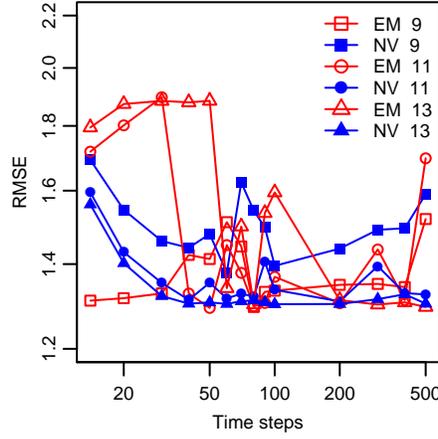
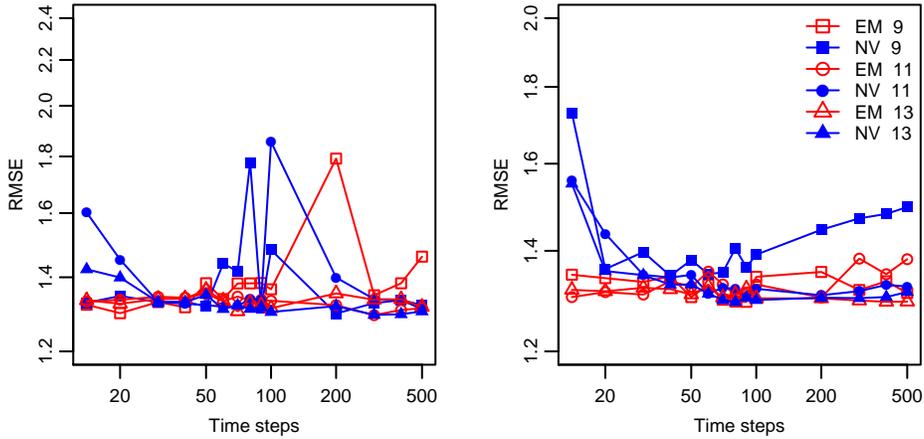


Figure 4: RMSE from NV and EM calibration of the DMR model to SPX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$, $\xi_1 = 2.873$, $\xi_2 = 0.302$ and the calibration is done with quasi random numbers using a Levenberg-Marquardt optimizer starting in $(-0.7, -0.7)$. The legend specifies the method and the \log_2 number of paths.



(a) Pseudo random numbers

(b) Quasi random numbers

Figure 5: RMSE from NV and EM calibration of the DMR model to SPX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$, $\xi_1 = 2.873$, $\xi_2 = 0.302$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (b) specifies the method and the \log_2 number of paths.

SPX option smiles very well, shorter expirations somewhat less well. This of course is not unexpected for a three-factor model of the DMR type. Note the kinks in the model implied volatilities for $T = 0.049$ and 0.20 . These appear

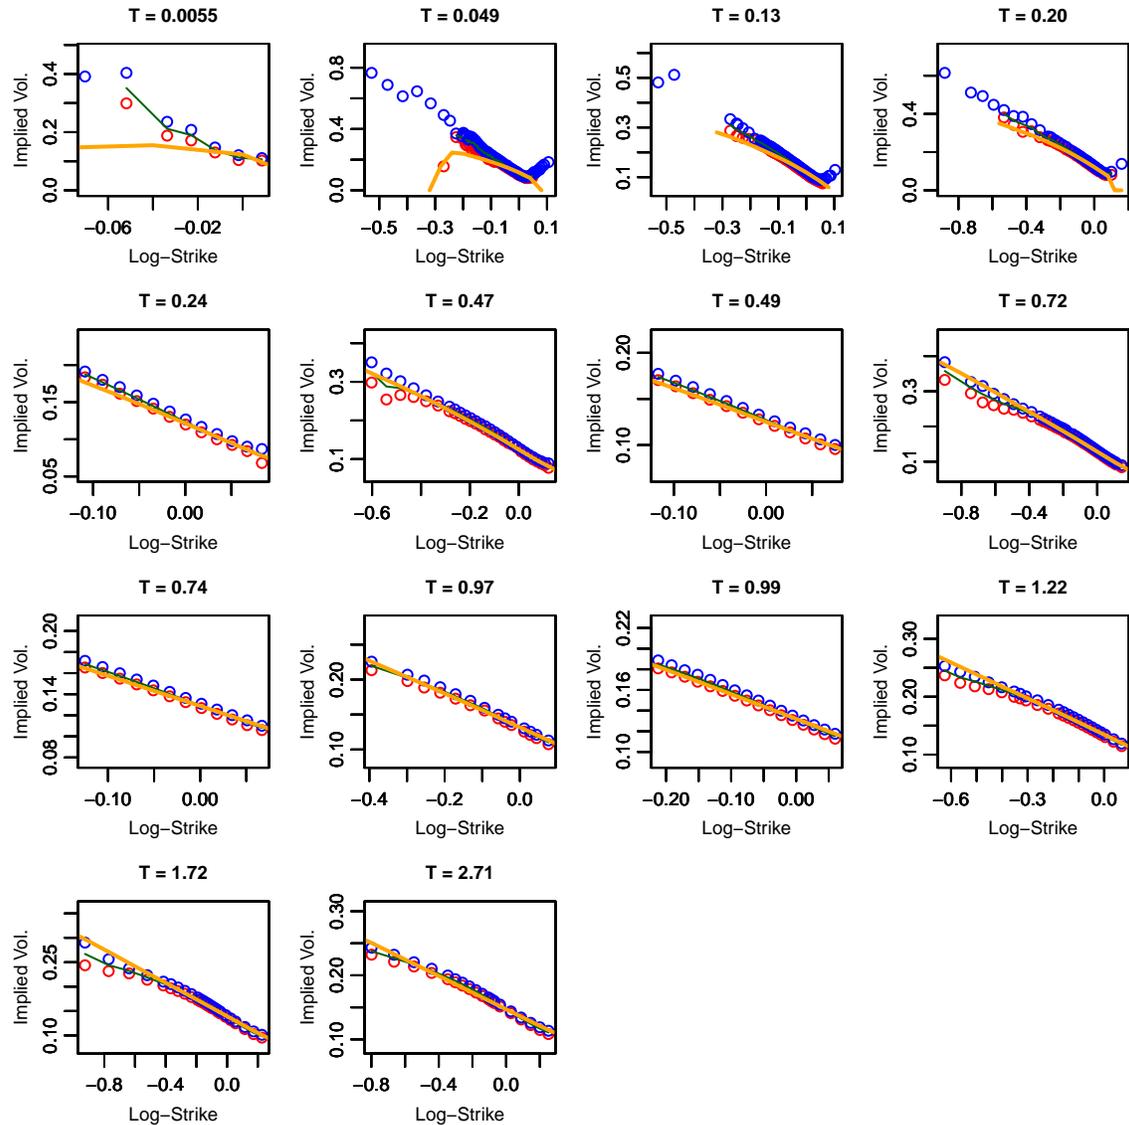


Figure 6: Implied Black volatilities for SPX options on April 3, 2007 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from a QMC-NV scheme using 100 time steps and 2^{16} paths (yellow line). The parameters for the model have been obtained using an NV scheme with 30 time steps and 2^{11} paths to calibrate the model to the VIX options, and an EM scheme with 30 time steps and 2^{11} paths to calibrate the model to the SPX options.

since we use too few paths, the model will therefore produce a time value of zero for options far from ATM. Finally, we observe that the calibrated parameters are quite consistent with those reported in Gatheral (2008), except ρ_{12} . However the calibration routine described here runs very much faster.

4.3.5 Calibration RMSE

One could argue that the preceding tests only examine whether or not the methods produce parameters that are good, and not that the schemes with the chosen number of paths and time steps fit the market prices. In figure 7 we graph the RMSE between the market prices and the model prices calculated with the same number of paths and timesteps that we use in the calibrations. We have for example obtained $\xi_1^{opt,30,11}$ and $\xi_1^{opt,30,11}$ the optimal parameters for a calibration with 30 time steps and 2^{11} paths. To obtain the calibration RMSE we first simulate the model using 30 time steps, 2^{11} paths, $\xi_1^{opt,30,11}$ and $\xi_1^{opt,30,11}$. Then we calculate the RMSE between these model prices and market prices.

For the VIX options we do not observe that big a difference from the previous RMSE graph. The RMSE only become a bit more bumpy, which is expected since we have higher integration error. The NV scheme with 2^{13} or 2^{11} paths can fit market prices with relatively few steps while the EM scheme needs at least 200 steps.

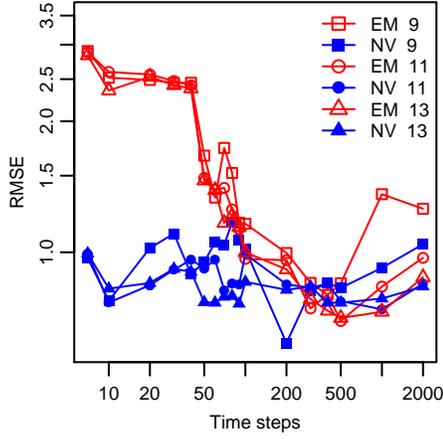
For SPX options we observe that the NV scheme with 2^{13} or 2^{11} paths and relatively few time steps generates good fits to market prices. In contrast, the EM scheme requires a large number of time steps to achieve the same fit quality. Nevertheless, if we are concerned only with calibration, the EM scheme does produce perfectly acceptable parameters with few time steps at significantly lower computational cost than NV.

4.4 September 15, 2011

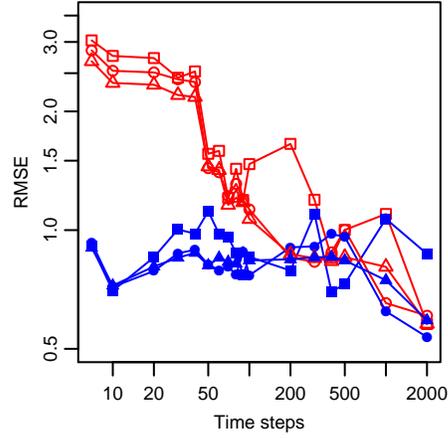
The dataset of Section 4.3 is from a period before the 2008 financial crisis. To investigate whether the model still works under more recent market conditions, and to further compare simulation schemes, we now calibrate the DMR model to data from September 15, 2011.

4.4.1 The data

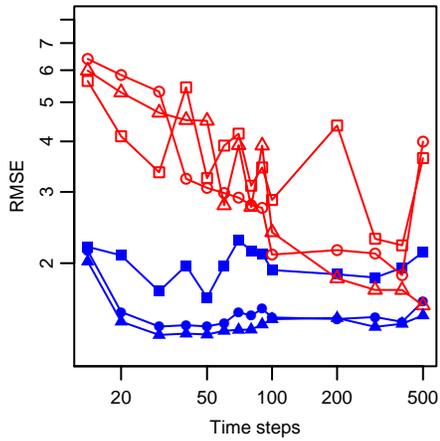
The SPX option dataset contains prices for 1176 options, 931 of them with both bid and ask prices, which we use in our calibration. The dataset contains 14 different maturities ranging from 0.0027 to 2.26 years. The strikes



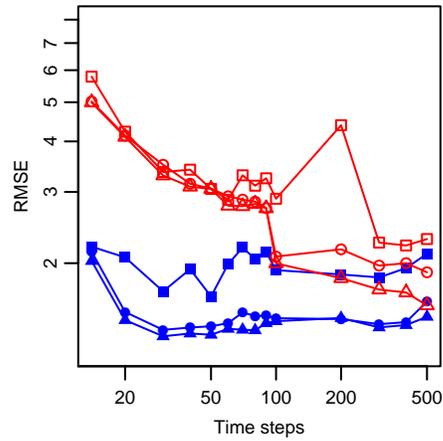
(a) Vix options, pseudo random numbers



(b) Vix options, quasi random numbers



(c) SPX options, quasi random numbers but no search algorithm



(d) SPX options, quasi random numbers and search algorithm

Figure 7: RMSE from NV and EM calibration of the DMR model to VIX and SPX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$ and for the calibration to SPX options we have also set $\xi_1 = 2.873$, $\xi_2 = 0.302$. Calibration to VIX options are considered in (a) with pseudo random numbers and in (b) with quasi random numbers. Calibration to SPX options are considered in (c) with quasi random numbers but without the search algorithm and in (d) with quasi random number and the search algorithm. The legend in (a) specifies the method and the \log_2 number of paths.

range from 100 to 3000. The forward for the first maturity is 1207.70 and for the last maturity 1159.83.

The VIX option data contains prices for 202 options, 148 of them have both bid and ask prices, again these are options we consider in our calibration. The dataset contains 6 different maturities ranging from 0.016 years to 0.42 years. The strikes ranges from 17 to 100. The forward is 33.23 for the first maturity and 31.29 for the last maturity.

4.4.2 Calibration of the parameters

Using linear regression to calibrate v_0 and v'_0 gives us the parameters

$$\begin{aligned} v_0 &= 0.114, \\ v'_0 &= 0.110. \end{aligned}$$

We then test how well the various simulation schemes calibrate the model to VIX options as in Section 4.3.3, presenting the results in figure 8.

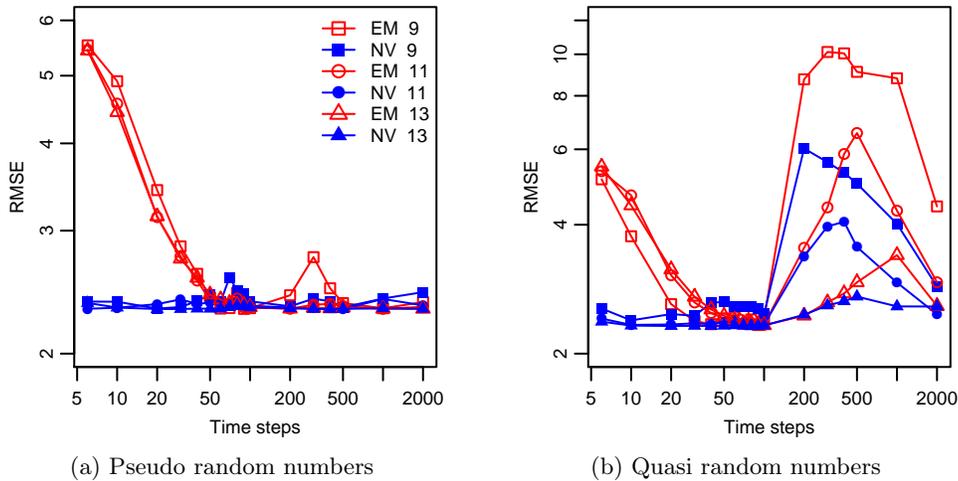


Figure 8: RMSE from NV and EM calibration of the DMR model to VIX option prices from the 15th of September 2011. We have set $\alpha_1 = \alpha_2 = 0.94$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (a) specifies the method and the \log_2 number of paths.

The NV discretization again beats the EM discretization but the difference is not as significant as in the 2007 example, requiring of the order of a factor 10 fewer time steps. The reasons for the decrease in relative performance are twofold: There are no long-dated VIX options in the 2011 dataset, and the volatility processes have higher starting values.

The longest VIX options have a maturity of 1.13 years in the 2007 dataset and 0.42 years in the 2011 dataset. We therefore expect the EM scheme to use $\frac{1.13}{0.42} = 2.69$ as many time steps in 2007 than in 2011, everything else being equal. Since the NV-scheme is second order convergent it will only use $\sqrt{2.69} = 1.64$ as many time steps in 2007 than in 2011.

The higher value for the volatility processes means that they will hit zero less frequently. This affects both schemes in a positive way. We can therefore reduce the number of timesteps for both schemes compared to the 2007 calibration. But there exists a lower bound, the number of VIX option maturities we need to hit. Therefore the number of timesteps cannot be lower than 6, even though the NV scheme could approximate the longest maturity well with fewer equidistant timesteps.

In Figure 9 we graph market VIX Black-Scholes implied volatility smiles together with model smiles. The calibrated ξ parameters are

$$\begin{aligned}\xi_1 &= 2.689, \\ \xi_2 &= 0.502.\end{aligned}$$

they were obtained using an MC-NV scheme with 6 time steps and 2^{11} paths. Model option prices were then computed using 100 NV QMC time steps and 2^{16} paths.

We observe that the DMR model generates VIX smiles that are too flat. This suggests that the lognormal DMR model with $\alpha_1 = \alpha_2 = 1$, which is faster to simulate, may also fit better. We investigate this in Section 4.5. The bid-ask spread have also decreased compared to the 2007 data, making it harder for the model to hit the market prices.

Using ξ_1 and ξ_2 just obtained we then calibrate correlation parameters. Results from this calibration are presented in Figure 10.

Again, we see no advantage in using the NV discretization over the simpler (and less costly) EM discretization. Again we conclude that quasi random numbers have to be used when calibrating to SPX options.

In Figure 12 we graph market SPX Black-Scholes implied volatility smiles together with model smiles. The model parameters ρ_{12} and ρ_{13} were calibrated to the market using QMC-EM scheme with 14 time steps, 2^{11} paths and search for a good starting point. The resulting calibrated ρ parameters are:

$$\begin{aligned}\rho_{12} &= -0.982, \\ \rho_{13} &= -0.727.\end{aligned}$$

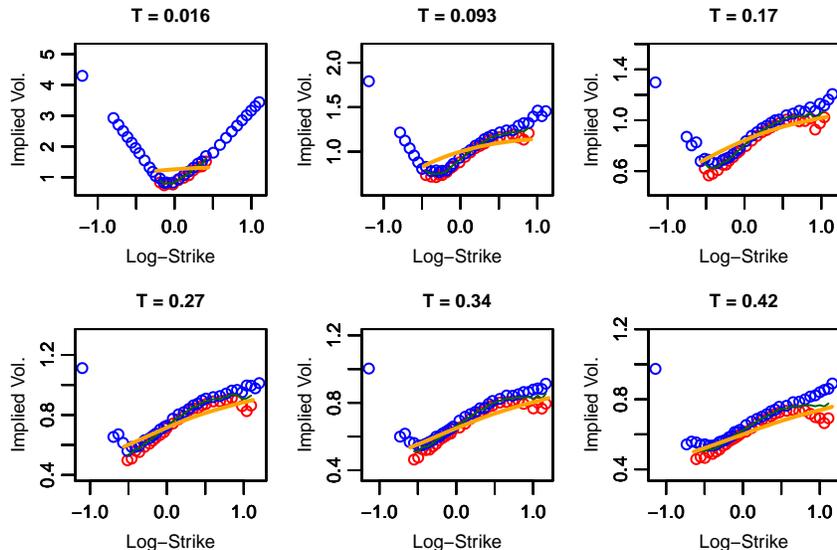


Figure 9: Implied Black volatilities for VIX options on September 15, 2011 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from a QMC-NV scheme using 100 time steps and 2^{16} paths (yellow line). The model parameters ξ_1 and ξ_2 are obtained by a calibration using a NV scheme with 6 time steps and 2^{11} Monte Carlo paths.

Model option prices were then computed using 100 NV QMC time steps and 2^{16} paths. As with the 2007 calibration, the DMR model fits SPX option prices well except for very short expirations.

4.5 Lognormal DMR model calibration to 2011 data

In the previous section we saw that a model where $\alpha_1 = \alpha_2 = 0.94$ generates VIX option smiles that are too flat compared to the market prices of our 2011 example. In order to increase the steepness of the smile we calibrate the simpler lognormal DMR model with $\alpha_1 = \alpha_2 = 1$.

In Figure 11 we graph the VIX option smiles obtained from the lognormal model. The volatility parameters have been obtained using an MC-NV scheme with 10 time steps and 2^{11} paths. The smiles have steepened compared to the graph in Section 4.4.2, the lognormal DMR model therefore fits the market better than the more complicated DMR model calibrated in Section 4.4.2. But the option smiles still seem to be too flat. From Figure 12 we see that calibrating the double lognormal model to the SPX options yields more or less the same smiles as before.

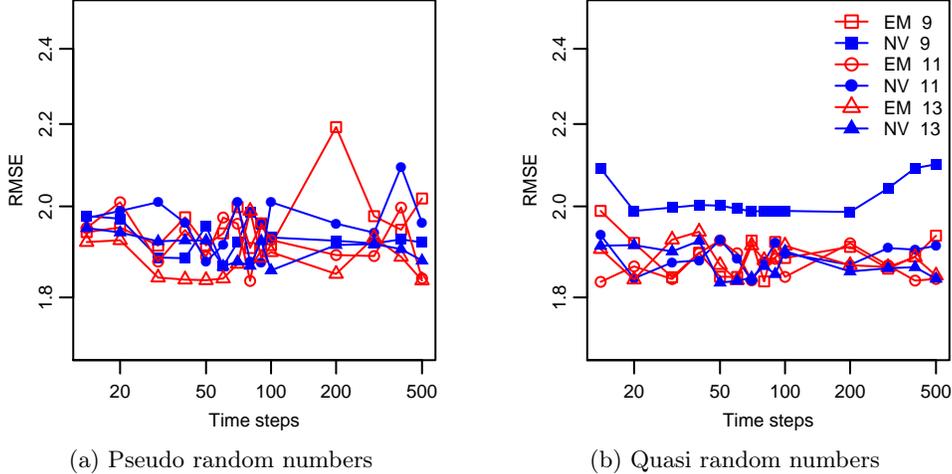


Figure 10: RMSE from NV and EM calibration of the DMR model to SPX option prices from the 15th of September 2011. We have set $\alpha_1 = \alpha_2 = 0.94$, $\xi_1 = 2.689$ and $\xi_2 = 0.502$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (b) specifies the method and the \log_2 number of paths.

4.6 Computation times

As we have seen, the NV-scheme can reduce by a large factor the number of time steps needed to achieve a good calibration of the DMR model to VIX options. In the 2007 example we got a reduction of a factor 15 – 20 and in the 2011 example we got a reduction of a factor 10. But this will not lead to an equivalent reduction in computation time because the NV discretization involves more computation and is therefore slower.

In Table 1, we present the empirical computational cost of the NV discretization relative to the EM discretization. The results are obtained using quasi random numbers, we get almost the same results if we use pseudo random numbers. Because of the drift trick, the NV discretization step is much simpler in the case $\alpha_1 = \alpha_2 = 1$ therefore we present this case separately. 2D is the case where we only simulate the variance processes i.e. when we have to price VIX options. 3D is the simulation of the full model.

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 4 times in the 2007 example, 2 in the 2011 example and 5 in the 2011 log-normal DMR example. In summary, the optimal calibration recipe appears to be:

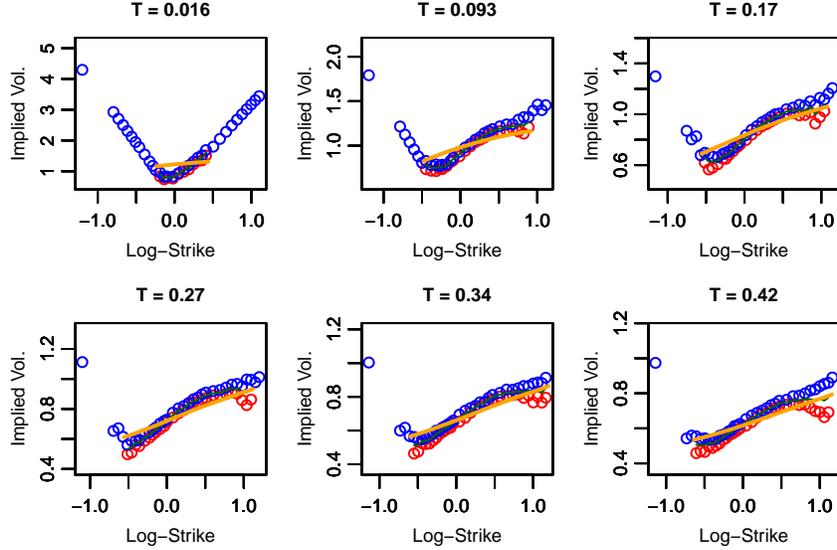


Figure 11: Implied Black volatilities for VIX option prices on September 15, 2011 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from QMC-NV with 100 time steps and 2^{16} paths (yellow line). The parameters are obtained by calibration using 10 NV time steps and 2^{11} QMC paths.

	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.

- 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 approximately 5 seconds.

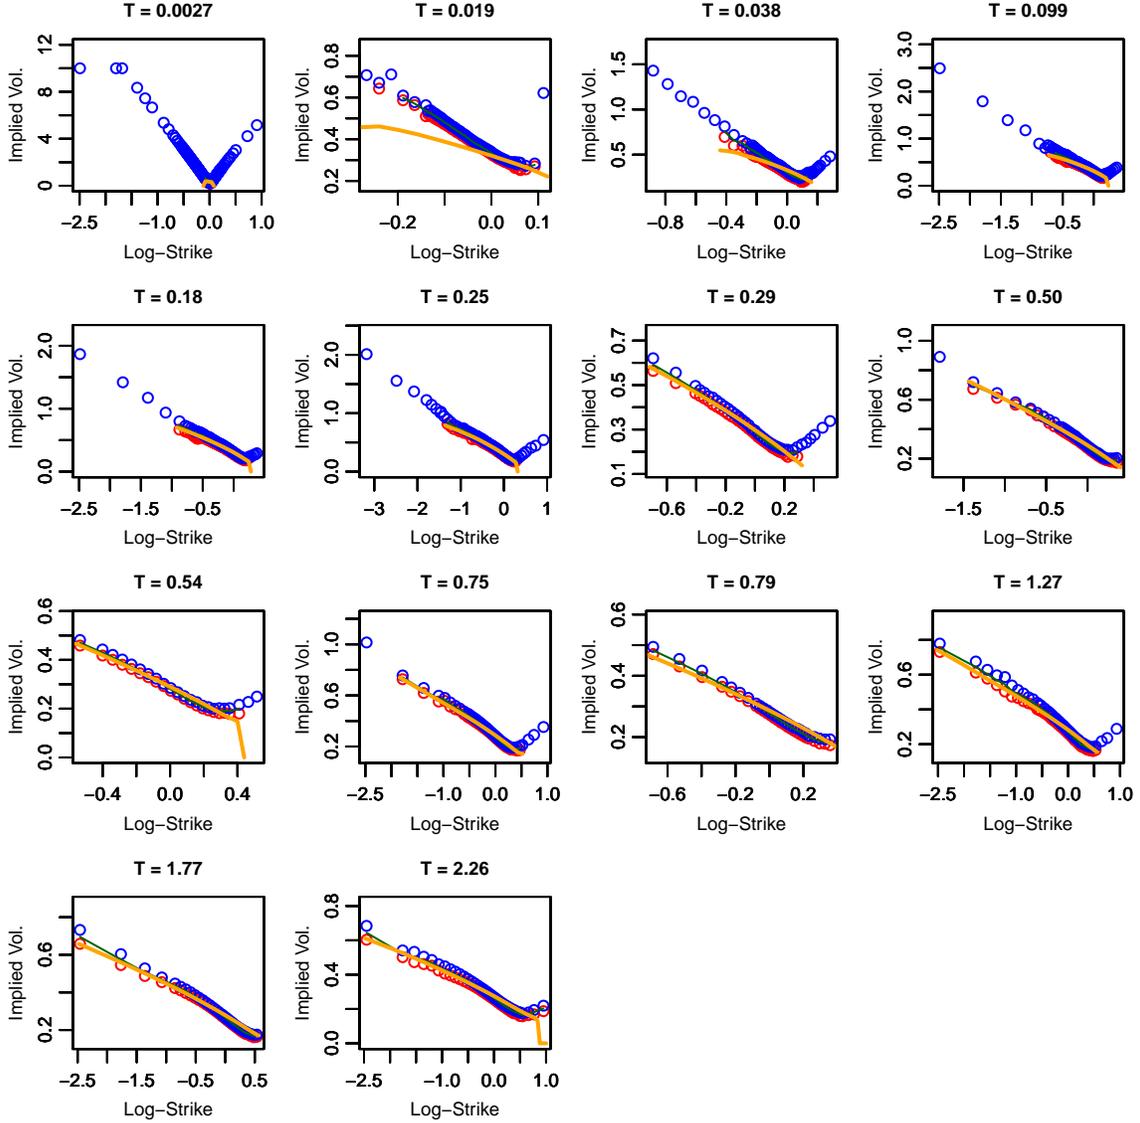


Figure 12: Implied Black volatilities for SPX options on September 15, 2011 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from the NV scheme using 100 time steps and 2^{16} Quasi Monte Carlo paths (yellow line). The parameters for the model have been obtained using the NV scheme with 6 time steps and 2^{11} Monte Carlo paths to calibrate the model to the VIX options, and an EM scheme with 14 time steps and 2^{11} Quasi Monte Carlo paths to calibrate the model to the SPX options.

5 Convergence of the discretization schemes

Having demonstrated in Section 4 that we have fast and accurate calibration of the DMR model to VIX and SPX options and moreover that fits to the market are good, we focus in this section on numerical tests of the convergence of the Ninomiya-Victoir (NV) discretization scheme presented in Section 3.2 relative to that of the Euler-Maruyama scheme with partial truncation presented in Section 2.3.

We use parameters resulting from the calibrations of Section 4.4 summarized in Table 2. We consider options with a maturity of one year and three

θ	0.078
κ_1	5.5
κ_2	0.1
ρ_{23}	0.59
v_0	0.114
v'_0	0.110
α_1	0.94
α_2	0.94
ξ_1	2.689
ξ_2	0.502
$\rho_{12} = \tilde{\rho}_{12}$	-0.982
$\rho_{13} = \tilde{\rho}_{13}$	-0.727
$\tilde{\rho}_{23}$	-0.656

Table 2: Parameters from the calibration to data from September 15, 2011 with $\alpha_1 = \alpha_2 = 0.94$.

different strikes, 0.8 times the forward, ATM and 1.2 times the forward. All option prices are computed using randomized QMC. The randomization is done by scrambling the net and then adding a random shift to the QMC numbers, see Glasserman (2004). This is done in order to obtain a Monte Carlo error around the price.

The "true" option prices are computed using 8 independent realizations, each realization is calculated using the NV scheme with 200 time steps and $2^{27} = 134,217,728$ QMC paths.

To obtain convergence graphs we calculate option prices using 5, 10, 20, 30, 50, 70 timesteps for the NV scheme and 5, 10, 20, 30, 50, 70, 100, 200, 500 timesteps for the EM scheme. We simulate 64 independent realizations of the NV prices and 128 independent realizations of the EM prices, each price is computed us-

ing $2^{23} = 8.388.608$ QMC paths. Convergence graphs can be seen in Figure 13.

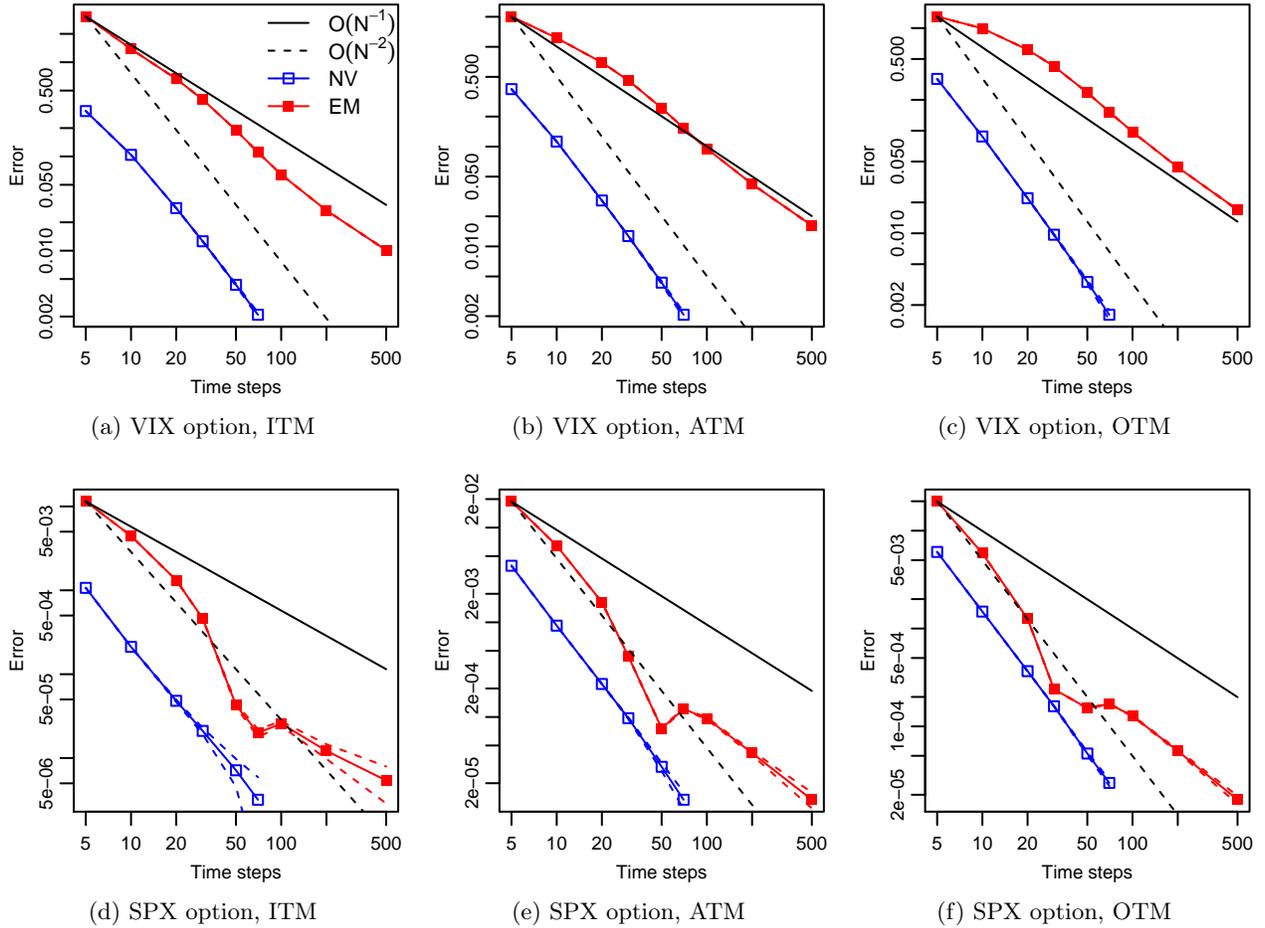


Figure 13: Pricing error for 1 year VIX and SPX options in a DMR model with parameters from 15th of September 2011 where $\alpha_1 = \alpha_2 = 0.94$. A confidence interval of two standard deviations around the error are marked with the dashed lines. The legend in (a) specifies the method.

The VIX option graphs in Figure 13 clearly show that the NV scheme has second order convergence while the EM scheme only converges with order one. For SPX options the picture is blurred a bit by the strange behavior of the EM scheme. There seems to be a kink in the error graph around 50 timesteps. The kink exists because the EM scheme with a small number of time steps creates too high option prices, while the EM scheme with a large

number of steps creates too low option prices. Therefore the EM scheme price has to cross the "true" price at one point, this point lies around 50 time steps. After the kink we see first order convergence of the EM scheme. The NV scheme clearly shows second order convergence.

In Lord et al. (2010) the Ninomiya-Victoir scheme was found inferior to the full truncation scheme when simulating the Heston model. In the Heston model the stochastic volatility is a CIR process ($\alpha_1 = 0.5$ and no v'_t), but in the DMR model we consider $\alpha_1 \approx 1$. If the parameters of the CIR process violate the Feller condition, 0 is an attainable boundary, and even if they satisfy the Feller condition the process can hit 0 when we simulate the model discretely. If α_1 is closer to one this will happen less frequently. We may therefore expect a method which is second order accurate for smooth volatility and drift functions with bounded derivatives to perform better when $\alpha_1 \approx 1$.

Figure 13 (d)-(f) clearly show that the Ninomiya-Victoir scheme comes much closer to the true price of SPX options for a given number of timesteps, this is also what we conclude from section 4.3.5. But we do not need very precise option prices to infer good correlation parameters. Given its significantly lower computational cost, the EM scheme is therefore to be preferred when calibrating to SPX options.

As for pricing VIX options, (which Lord et al. (2010) do not consider), our tests show that the outperformance of the NV scheme is sufficient for it to be preferred over the EM scheme for the calibration of volatility parameters.

6 Conclusion

In this paper, 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, avoiding the use of numerical integration methods such as Runge-Kutta. 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 of Gatheral (2008) 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.

Finally we performed an empirical analysis of the convergence of the

Ninomiya-Victoir (NV) and Euler-Maruyama discretization schemes demonstrating that the NV scheme was indeed second-order weak convergent.

References

- Alfonsi, A. (2010). High order discretization schemes for the CIR process: application to affine term structure and Heston models. *Math. Comp.* 79(269), 209–237.
- Andersen, T., T. Bollerslev, F. Diebold, and H. Ebens (2001). The distribution of realized stock return volatility. *Journal of Financial Economics* 61(1), 43–76.
- Bayer, C., P. Friz, and R. Loeffen (2013). Semi-closed form cubature and applications to financial diffusion models. *Quantitative Finance*. To appear.
- Bergomi, L. (October 2005). Smile dynamics II. *Risk* 18, 67–73.
- Buehler, H. (2006). Consistent variance curve models. *Finance and Stochastics* 10, 178–203.
- da Silva, M. E. and T. Barbe (2005). Quasi-Monte Carlo in finance: Extending for problems of high effective dimension. *Economia Aplicada* 9(4).
- Gatheral, J. (2006). *The Volatility Surface: A Practitioner’s Guide*. Hoboken, New Jersey: John Wiley & Sons, Inc.
- Gatheral, J. (2008). Consistent modelling of SPX and VIX options. Bachelier congress, London.
- Gatheral, J. and A. Jacquier (2012). Arbitrage-free SVI volatility surfaces. SSRN.
- Glasserman, P. (2004). *Monte Carlo Methods In Financial Engineering*. New York, New York: Springer.
- Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002). Managing smile risk. *Wilmott Magazine* 1(1), 84–108.
- Hairer, E., C. Lubich, and G. Wanner (2006). *Geometric numerical integration* (Second ed.), Volume 31 of *Springer Series in Computational Mathematics*. Berlin: Springer-Verlag. Structure-preserving algorithms for ordinary differential equations.
- Joe, S. and F. Y. Kou (2008). Constructing sobol sequences with better two-dimensional projections. *SIAM J. Sci. Comput* 30, 2635–2654.

- Karatzas, I. and S. Shreve (1988). *Brownian Motion and Stochastic Calculus*. New York, Heidelberg & Berlin: Springer-Verlag.
- Lord, R., R. Koekkoek, and D. van Dijk (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance* 10(2), 177–194.
- Lyons, T. and N. Victoir (2004). Cubature on Wiener space. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* 460(2041), 169–198. Stochastic analysis with applications to mathematical finance.
- More, J. J., B. S. Garbow, and K. E. Hillstrom (1980). User guide for minpack-1. Technical report, Argonne National Laboratory.
- Ninomiya, M. and S. Ninomiya (2009). A new higher-order weak approximation scheme for stochastic differential equations and the Runge-Kutta method. *Finance Stoch.* 13(3), 415–443.
- Ninomiya, S. and N. Victoir (2008). Weak approximation of stochastic differential equations and application to derivative pricing. *Applied Mathematical Finance* 15, 107–121.
- Strang, G. (1963). Accurate partial difference methods. I. Linear Cauchy problems. *Arch. Rational Mech. Anal.* 12, 392–402.
- Talay, D. and L. Tubaro (1990). Expansion of the global error for numerical schemes solving stochastic differential equations. *Stochastic Anal. Appl.* 8(4), 483–509 (1991).