



HAL
open science

The quintic Ornstein-Uhlenbeck volatility model that jointly calibrates SPX & VIX smiles

Eduardo Abi Jaber, Camille Illand, Shaun Xiaoyuan Li

► **To cite this version:**

Eduardo Abi Jaber, Camille Illand, Shaun Xiaoyuan Li. The quintic Ornstein-Uhlenbeck volatility model that jointly calibrates SPX & VIX smiles. 2022. hal-03909334v1

HAL Id: hal-03909334

<https://hal.science/hal-03909334v1>

Preprint submitted on 21 Dec 2022 (v1), last revised 9 May 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The quintic Ornstein-Uhlenbeck volatility model that jointly calibrates SPX & VIX smiles

Eduardo Abi Jaber^{*1}, Camille Illand², and Shaun (Xiaoyuan) Li^{†3}

¹Ecole Polytechnique, CMAP

^{2,3}AXA Investment Managers

³Université Paris 1 Panthéon-Sorbonne, CES

December 21, 2022

Abstract

The quintic Ornstein-Uhlenbeck volatility model is a stochastic volatility model where the volatility process is a polynomial function of degree five of a single Ornstein-Uhlenbeck process with fast mean reversion and large vol-of-vol. The model is able to achieve remarkable joint fits of the SPX-VIX smiles with only 6 effective parameters and an input curve that allows to match certain term structures. Even better, the model remains very simple and tractable for pricing and calibration: the VIX squared is again polynomial in the Ornstein-Uhlenbeck process, leading to efficient VIX derivative pricing by a simple integration against a Gaussian density; simulation of the volatility process is exact; and pricing SPX products can be done efficiently and accurately by standard Monte Carlo techniques with suitable antithetic and control variates.

JEL Classification: G13, C63, G10.

Keywords: SPX and VIX modeling, Stochastic volatility, Pricing, Calibration.

1 Introduction

Since the financial crisis of 2008, derivatives on volatility became increasingly popular for hedging purposes and for directional trading especially when combined with the underlying stock index. On the US market, the VIX index introduced by the CBOE became one of the most widely followed volatility index. By construction, the VIX expresses an interpolation between several points of the SPX implied volatility term structure. This motivates the need for a consistent modeling of the SPX and VIX.

By joint SPX-VIX calibration problem, we mean the calibration of a model across several maturities to European call/put options on SPX and VIX together with VIX futures. Such joint calibration turns out to be quite challenging for several reasons: multitude of instruments to be calibrated (SPX/VIX options and VIX futures, so three types of derivatives) across several maturities (to stay consistent with the construction of the VIX), characterized by an upward sloping VIX

^{*}eduardo.abi-jaber@polytechnique.edu. The first author is grateful for the financial support from the Chaires FiME-FDD, Financial Risks, Deep Finance & Statistics and Machine Learning and systematic methods in finance at Ecole Polytechnique.

[†]shaunlinz02@gmail.com.

implied volatility in contrast with the important at-the-money (ATM) SPX skew that becomes more pronounced for smaller maturities.

Several attempts of joint calibration have been made, some of which are rather satisfactory. However, in general the models and/or the techniques considered are sophisticated and make use of jump processes [3, 6, 17, 20, 21]; non-Markovian rough volatility [4, 10, 23] and path-dependent volatility [16]; multiple-factors [7, 11, 16, 22]; optimal transport [13, 14, 15] and randomization of the parameters [12] to just name a few. The proposed solutions for the joint calibration problem are hence rather challenging to put into practice and need specific advanced numerical methods. This is the main motivation of our work.

Recently, the work of [2] identified for the first time a conventional one-factor Markovian continuous stochastic volatility model that is capable of achieving remarkable fits for a wide range of maturity slices of SPX and VIX implied volatilities together with the term structure of VIX futures. It is also shown on an extensive empirical study between 2012 and 2022, that contrary to common beliefs, the one factor Markovian model can jointly calibrate SPX and VIX without appealing to multiple-factors, jumps, rough volatility or path-dependency and can achieve even better performances. In [2], pricing of VIX and SPX derivatives has been done using quantization techniques and neural networks in order to ensure a fair comparison between Markovian and non-Markovian models.

In the present work, we focus on the Markovian model identified in [2] and we show that the model is tractable in addition to being remarkably flexible. The dynamics of the stochastic volatility process in this model are given by a polynomial function of degree five of a single Ornstein-Uhlenbeck process with fast mean reversion and large vol-of-vol. Hence the name: quintic Ornstein-Uhlenbeck volatility model. The model has only 6 effective parameters and an input curve that allows to match certain term structures. In particular, we will highlight the role of the input curve on the joint calibration problem and consider additionally time-dependent parameters to help achieve better fits for longer maturities.

We show that the model is tractable as it offers an explicit expression for the VIX squared which is again polynomial in the driving Ornstein-Uhlenbeck factor, leading to efficient VIX derivative pricing by integrating directly against a Gaussian density. Simulation of the volatility process is exact so that pricing SPX products can be done efficiently and accurately by standard Monte Carlo techniques with suitable antithetic and control variates. We also provide a notebook with our implementation here: https://colab.research.google.com/drive/14nh9civ_wgQv283eshBwnr146w7Xsbi5?usp=sharing.

For the first time in the literature, remarkable joint fits of SPX and VIX volatility surfaces and VIX futures are achieved between 1 week and beyond 1 year. Although it is challenging, but possible, for another model to achieve similar fits, it would be very difficult to do so with a simpler continuous model than our quintic Ornstein-Uhlenbeck volatility model.

Outline. Section 2 introduces the model. Sections 3 and 4 detail the pricing of VIX and SPX derivatives in the model. Calibration results on market data are shown in Section 5.

2 The quintic Ornstein-Uhlenbeck volatility model

The dynamics of the stock price S , with no interest nor dividends, is given by

$$\begin{aligned} \frac{dS_t}{S_t} &= \sigma_t dB_t, \\ \sigma_t &= \sqrt{\xi_0(t)} \frac{p(X_t)}{\sqrt{\mathbb{E}[p(X_t)^2]}}, \quad p(x) = \alpha_0 + \alpha_1 x + \alpha_3 x^3 + \alpha_5 x^5, \\ X_t &= \varepsilon^{H-1/2} \int_0^t e^{-(1/2-H)\varepsilon^{-1}(t-s)} dW_s, \end{aligned} \tag{2.1}$$

with $B = \rho W + \sqrt{1 - \rho^2} W^\perp$, (W, W^\perp) a two-dimensional Brownian motion on a risk-neutral filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$, $\rho \in [-1, 1]$, non-negative coefficients $\alpha_0, \alpha_1, \alpha_3, \alpha_5 \geq 0$ ($\alpha_2 = \alpha_4 = 0$), $\varepsilon > 0$, $H \in (\infty, 1/2]$ and an input curve $\xi_0 \in L^2([0, T], \mathbb{R}_+)$ for any $T > 0$, allowing the model to match certain term-structures observed on the market. For instance, the normalization $\sqrt{\mathbb{E}[p(X_t)^2]}$ allows ξ_0 to match the market forward variance curve since

$$\mathbb{E} \left[\int_0^t \sigma_s^2 ds \right] = \int_0^t \xi_0(s) ds, \quad t \geq 0.$$

The process X driving the volatility is an Ornstein-Uhlenbeck process with a fast mean reversion of order $(1/2 - H)\varepsilon^{-1}$ and a large vol-of-vol of order $\varepsilon^{H-1/2}$ for small values of ε , that is

$$dX_t = -(1/2 - H)\varepsilon^{-1} X_t dt + \varepsilon^{H-1/2} dW_t. \quad (2.2)$$

Such parametrizations are reminiscent of the fast regimes extensively studied by Fouque et al. [9], see also [8, Section 3.6], which corresponds to the case $H = 0$. They can also be linked to more complex models such as jump models [19, 1] for $H \leq -1/2$; and rough volatility models [2, 1] for which $H \in (0, 1/2)$ would play the role of the Hurst index. Letting the parameter $H \in (\infty, 1/2]$ free in our model introduces more flexibility and leads to better fits than in the aforementioned models. Another advantage of such parametrization is to stabilize the calibrated value of H through time as opposed to calibrating directly on mean reversion and vol-of-vol parameters which are less stable through time, see [2, Figure 3].

Taking p a polynomial of degree five allows us to reproduce the upward slope of the VIX smile. Restricting the coefficients α to be non-negative allows the sign of the at-the-money skew to be the same as ρ , see [2] for more details. Setting $\alpha_2 = \alpha_4 = 0$ allows us to reduce the number of parameters to calibrate.

We fix $\varepsilon = 1/52$ to further reduce the parameters, which gives 6 calibratable parameters:

$$\Theta := \{\alpha_0, \alpha_1, \alpha_3, \alpha_5, \rho, H\}, \quad (2.3)$$

plus the input curve $\xi_0(\cdot)$. Numerical experiments show no significant adverse impact on the joint calibration quality by narrowing the choice of parameters.

3 Pricing VIX derivatives

An explicit expression for the VIX. One major advantage of our model is an explicit expression of the VIX. In continuous time, the VIX can be expressed as

$$\text{VIX}_T^2 = -\frac{2}{\Delta} \mathbb{E}[\log(S_{T+\Delta}/S_T) \mid \mathcal{F}_T] \times 100^2 = \frac{100^2}{\Delta} \int_T^{T+\Delta} \xi_T(u) du, \quad (3.1)$$

with $\Delta = 30$ days and $\xi_T(u) := \mathbb{E}[\sigma_u^2 \mid \mathcal{F}_T]$ the forward variance process which can be computed explicitly in our model as follows. First, we fix $T \leq u$ and rewrite X as

$$X_u = X_T e^{-(1/2-H)\varepsilon^{-1}(u-T)} + \varepsilon^{H-1/2} \int_T^u e^{-(1/2-H)\varepsilon^{-1}(u-s)} dW_s =: Z_T^u + G_T^u,$$

then, setting

$$g(u) = \mathbb{E}[p(X_u)^2],$$

we have that

$$\xi_T(u) = \mathbb{E}[\sigma_u^2 \mid \mathcal{F}_T] = \frac{\xi_0(u)}{g(u)} \mathbb{E} \left[\left(\sum_{k=0}^5 \alpha_k X_u^k \right)^2 \mid \mathcal{F}_T \right] = \frac{\xi_0(u)}{g(u)} \mathbb{E} \left[\sum_{k=0}^{10} (\alpha * \alpha)_k X_u^k \mid \mathcal{F}_T \right],$$

where $(\alpha * \alpha)_k = \sum_{j=0}^k \alpha_j \alpha_{k-j}$ is the discrete convolution. Using the Binomial expansion, we can further develop the expression for $\xi_T(u)$ in terms of Z^u and G^u to get

$$\xi_T(u) = \frac{\xi_0(u)}{g(u)} \sum_{k=0}^{10} \sum_{i=0}^k (\alpha * \alpha)_k \binom{k}{i} \left(X_T e^{-(1/2-H)\varepsilon^{-1}(u-T)} \right)^i \mathbb{E} [(G_T^u)^{k-i}], \quad (3.2)$$

where we used the fact that Z_T^u is \mathcal{F}_T -measurable and that G_T^u is independent of \mathcal{F}_T , with $\binom{k}{i} = k! / ((k-i)! i!)$ the binomial coefficient. Furthermore, G_T^u is a Gaussian random variable with mean 0 and variance $\frac{\varepsilon^{2H}}{1-2H}(1 - e^{-(1-2H)\varepsilon^{-1}(u-T)})$. Recall that for a Gaussian variable $Y \sim \mathcal{N}(0, \sigma_Y^2)$, its moments $\mathbb{E}[Y^p]$ for $p \in \mathbb{N}$ can be computed explicitly:

$$\mathbb{E}[Y^p] = \begin{cases} 0 & \text{if } p \text{ is odd} \\ \sigma_Y^p (p-1)!! & \text{if } p \text{ is even} \end{cases} \quad (3.3)$$

with $p!!$ the double factorial. Therefore all moments of $\mathbb{E}[(G_T^u)^i]$ are given explicitly.

Going back to (3.1) and plugging the expression (3.2), the explicit expression of the VIX_T^2 turns out to be polynomial in X_T :

$$\begin{aligned} \text{VIX}_T^2 &= \frac{100^2}{\Delta} \sum_{k=0}^{10} \sum_{i=0}^k (\alpha * \alpha)_k \binom{k}{i} \int_T^{T+\Delta} \frac{\xi_0(u)}{g(u)} \mathbb{E} [(G_T^u)^{k-i}] e^{-(1/2-H)\varepsilon^{-1}(u-T)i} du X_T^i \\ &= \frac{100^2}{\Delta} \sum_{i=0}^{10} \sum_{k=i}^{10} \left((\alpha * \alpha)_k \binom{k}{i} \int_T^{T+\Delta} \frac{\xi_0(u)}{g(u)} \mathbb{E} [(G_T^u)^{k-i}] e^{-(1/2-H)\varepsilon^{-1}(u-T)i} du \right) X_T^i \\ &= \frac{100^2}{\Delta} \sum_{i=0}^{10} \beta_i X_T^i, \end{aligned} \quad (3.4)$$

where

$$\beta_i = \sum_{k=i}^{10} (\alpha * \alpha)_k \binom{k}{i} \int_T^{T+\Delta} \frac{\xi_0(u)}{g(u)} \mathbb{E} [(G_T^u)^{k-i}] \left(e^{-(1/2-H)\varepsilon^{-1}(u-T)i} \right) du.$$

Recall that the moments $\mathbb{E}[(G_T^u)^i]$ are given explicitly by (3.3). The integral inside β_i can be easily computed, at least numerically for a variety of choices for $\xi_0(\cdot)$.

Pricing VIX derivatives. Thanks to the closed expression of (3.4), VIX_T^2 is a polynomial in X_T that we denote by $h(X_T)$. Since X_T is Gaussian with mean 0 and variance $\sigma_{X_T}^2 = \frac{\varepsilon^{2H}}{1-2H}(1 - e^{-(1-2H)\varepsilon^{-1}T})$, pricing VIX derivatives with payoff function Φ is immediate by integrating directly against the standard Gaussian density:

$$\mathbb{E}[\Phi(\text{VIX}_T)] = \mathbb{E} \left[\Phi \left(\sqrt{h(X_T)} \right) \right] = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} \Phi \left(\sqrt{h(\sigma_{X_T} x)} \right) e^{-x^2/2} dx. \quad (3.5)$$

Example 3.1. To price VIX future prices, set $\Phi(v) = v$ and to price VIX vanilla call price, set $\Phi(v) = (v - K)^+$. This integral (3.5) can be computed efficiently using a variety of quadrature techniques. The Gaussian quadrature with 400 nodes seems to be more than enough to price accurately VIX call and future prices.

4 Pricing SPX derivatives

To price SPX derivatives, we resort to using Monte Carlo simulations. Since X is a Ornstein-Uhlenbeck process, it can be simulated exactly as opposed to using the Euler scheme which is

often inaccurate in a fast mean reversion regime. To simulate X , first define

$$\tilde{X}_t = X_t e^{\frac{1/2-H}{\varepsilon}t} = \varepsilon^{H-1/2} \int_0^t e^{\frac{1/2-H}{\varepsilon}s} dW_s.$$

Then, \tilde{X} can be simulated recursively by

$$\tilde{X}_{t_{i+1}} = \tilde{X}_{t_i} + \sqrt{\varepsilon^{2H}/(1-2H)} \left(e^{\frac{1-2H}{\varepsilon}t_{i+1}} - e^{\frac{1-2H}{\varepsilon}t_i} \right) Y_i,$$

with Y_i i.i.d. standard Gaussian. To get back to $X_{t_{i+1}}$ we just divide $\tilde{X}_{t_{i+1}}$ by $e^{\frac{1/2-H}{\varepsilon}t_{i+1}}$. This setting allows us to easily vectorize computations.

To simulate the process $\log(S)$, we use the Euler scheme together with antithetic and control variates, the so called turbocharging method as outlined in [18]. This means we only need to simulate the part of $\log(S)$ that is \mathcal{F}^W measurable, we call this S^W and can be simulated as

$$\log(S^W)_{t_{i+1}} = \log(S^W)_{t_i} - 1/2 (\rho\sigma_{t_i})^2 (t_{i+1} - t_i) + \rho\sigma_{t_i} \sqrt{t_{i+1} - t_i} Y_i.$$

The main idea of the turbocharging method is to 1) take advantage of the conditional log-normality of S with respect to \mathcal{F}^W , hence removing the MC error from simulating W^\perp , and 2) apply the control variate in the form of a time option where one can again take advantage of the log-normality and closed form solution. We refer readers to [18] for more details on the method and to the notebook mentioned in the introduction for our implementation.

5 SPX/VIX Joint calibration

We now address the SPX-VIX joint calibration problem, that is the calibration of our model to SPX European options, VIX European options and VIX futures across several maturities. Ideally, one should calibrate for SPX options maturity up to one month ahead of that of the VIX options, given that VIX encodes expected level of volatility for the next 30 days by definition.

The calibration of VIX futures is necessary as it is used to calculate VIX implied volatility. Recall the implied volatility is calculated by inverting the Black and Scholes formula, that is, for a given call price $C_0(K, T)$ with strike K and maturity T , we find the unique $\sigma(K, T)$ such that

$$C_0(K, T) = F(T)\mathcal{N}(d_1) - K\mathcal{N}(d_2)$$

with

$$d_1 = \frac{\log(F(T)/K) + \frac{1}{2}\sigma(K, T)^2 T}{\sigma(K, T)\sqrt{T}}, \quad d_2 = d_1 - \sigma(K, T)\sqrt{T},$$

where $\mathcal{N}(x)$ is the cumulative density function of the standard Gaussian distribution and $F(T)$ denotes the futures price of the index: $F(T) = \mathbb{E}[S_T] = S_0$ for the SPX in our model (2.1) and $F(T) = \mathbb{E}[\text{VIX}_T]$ for the VIX.

To calibrate our model, we solve the following optimisation problem involving sum of root mean squared error (RMSE):

$$\min_{\Theta} \left\{ c_1 \sqrt{\sum_{i,j} \left(\sigma_{spx}^\Theta(T_i, K_j) - \sigma_{spx}^{mkt}(T_i, K_j) \right)^2} + c_2 \sqrt{\sum_{i,j} \left(\sigma_{vix}^\Theta(T_i, K_j) - \sigma_{vix}^{mkt}(T_i, K_j) \right)^2} \right. \\ \left. + c_3 \sqrt{\sum_i \left(F_{vix}^\Theta(T_i) - F_{vix}^{mkt}(T_i) \right)^2} \right\}.$$

Here, $\sigma_{spx}^{mkt}(T_i, K_j)$, $\sigma_{vix}^{mkt}(T_i, K_j)$ represent market SPX-VIX implied volatility with maturity T_i and strike K_j . $F_{vix}^{mkt}(T_i)$ is the market VIX futures price maturing at T_i . $\sigma_{spx}^\Theta(T_i, K_j)$, $\sigma_{vix}^\Theta(T_i, K_j)$

and $F_{vix}^\Theta(T_i)$ represent the same instruments, but coming from our model. The coefficients c_1 , c_2 and c_3 are some positive numbers used to assign different weights to the errors in SPX-VIX implied volatility and VIX futures price. We chose arbitrarily $c_1 = 1$, $c_2 = 0.1$, $c_3 = 0.5$ for our numerical experiments.

We will now show how our model is well adapted to produce joint fits between SPX/VIX, with daily SPX/VIX joint implied volatility surface data purchased from the CBOE website <https://datashop.cboe.com/>.

Extracting the forward variance curve $\xi_0(\cdot)$. Using the well-known replication formula for the log-contract in [5], we construct $\xi_0(\cdot)$ such that

$$\int_{T_i}^{T_{i+1}} \xi_0(s) ds = 2 \left(\int_0^{S_0} \frac{P_0(K, T_{i+1})}{K^2} dK + \int_{S_0}^{+\infty} \frac{C_0(K, T_{i+1})}{K^2} dK \right) - 2 \left(\int_0^{S_0} \frac{P_0(K, T_i)}{K^2} dK + \int_{S_0}^{+\infty} \frac{C_0(K, T_i)}{K^2} dK \right), \quad (5.1)$$

where T_i are SPX option maturities from market data and $C_0(K, T)$ and $P_0(K, T)$ the price of a call/put option with strike K and maturity T . Since market prices for out of the money call/put options are not always available, we first interpolate the SPX market implied volatility surface, each slice separately, using methods like SVI or SABR (after checking for arbitrage), and use the fitted surface to compute the integral above.

We then approximate $\xi_0(t)$ by passing a cubic spline interpolation with nodes (t_i, x_i) with $t_i = (T_i + T_{i+1})/2$ and $x_i = \sqrt{\int_{T_i}^{T_{i+1}} \xi_0(s) ds}$ and then square the interpolation to ensure positivity of the forward variance curve. Of course, piece-wise constant between $[T_i, T_{i+1})$ can also be used.

During the calibration procedure, we will let the optimisation algorithm move the model parameters Θ defined in (2.3) and make adjustments to the value of the spline nodes x_i as necessary to jointly fit SPX and VIX derivatives.

Figure 1 shows the joint fit on the 23 October 2017, with calibrated parameters $\rho = -0.6843$, $H = -0.0358$, $(\alpha_0, \alpha_1, \alpha_3, \alpha_5) = (0.5907, 1, 0.2893, 0.0549)$:

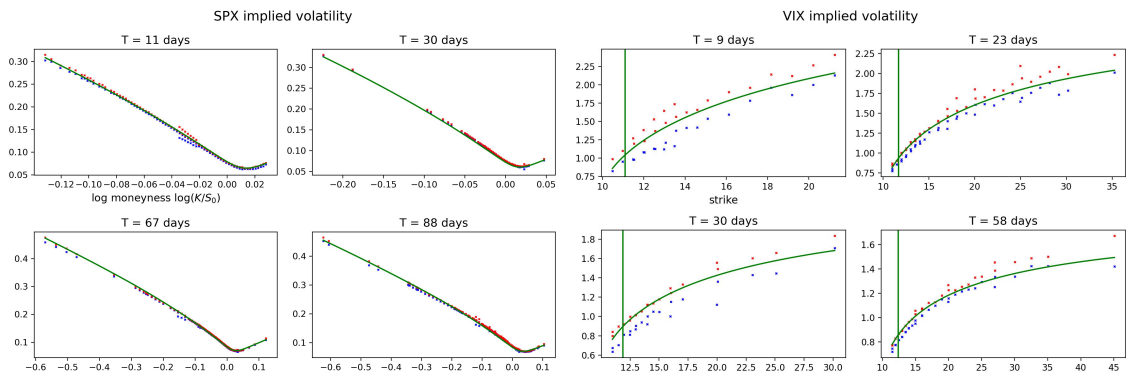


Figure 1: SPX-VIX smiles (bid/ask in blue/red) and VIX futures (vertical black lines) jointly calibrated with our model (full green lines) for 23 October 2017.

The forward variance curve has been adjusted to jointly fit the SPX and VIX smiles as shown in Figure 2.




Figure 2: The blue line represent the forward variance curve stripped from market data using Carr-Madan log contract formula in (5.1), the dotted green line is the adjusted forward variance curve as part of the calibration to jointly fit SPX/VIX smiles on 23 October 2017, with the round points representing cubic spline nodes.

For more joint surface fits and an empirical study on joint SPX/VIX volatility surface between 2011 and 2022 for our quintic Ornstein-Uhlenbeck model together with its calibrated parameters across time, we refer the reader to [2].

Using parametric forward variance curves when calibrating fewer slices. Instead of extracting forward variance curves from market data, it is also possible to use a parametric form of the forward variance curve for example in the form of:

$$\xi_0(t) = ae^{-bt} + c(1 - e^{-bt}), \quad (5.2)$$

with $a, b, c > 0$ to be calibrated.

The parametric forward variance curves offers less flexibility than that of extracted market forward variance curve discussed before given its rigid form. However, it is still capable to fit two maturity slices of SPX and one slice of VIX. We provide two examples here, with

1. joint fits of SPX options maturing in 9 days and 30 days, and VIX options maturing in 9 days) using parameters $\rho = -0.7316, H = -0.1382, (\alpha_0, \alpha_1, \alpha_3, \alpha_5) = (0.8169, 0.274, 0.1717, 0.0036), a = 0.0084, b = 2.0436, c = 0.0441$ shown in Figure 3,
2. joint fits of SPX options maturing in 53 days and 88 days, and VIX options maturing in 58 days) with parameters $\rho = -0.7001, H = 0.141, (\alpha_0, \alpha_1, \alpha_3, \alpha_5) = (0.7558, 1, 0.0885, 0.4421), a = 0.012, b = 2.027, c = 0.033$ shown in Figure 4.

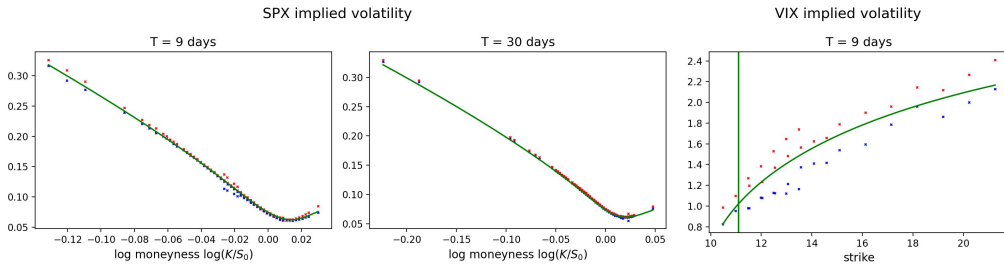


Figure 3: SPX-VIX smiles (bid/ask in blue/red) and VIX futures (vertical black lines) jointly calibrated with our model with the parametric forward variance curve (5.2) for 23 October 2017.

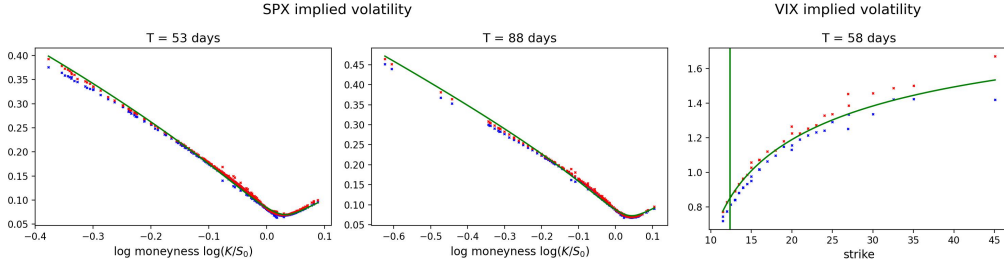


Figure 4: SPX–VIX smiles (bid/ask in blue/red) and VIX futures (vertical black lines) jointly calibrated with our model with the parametric forward variance curve (5.2) for 23 October 2017.

Using a time dependent H parameter for fitting longer maturities. To fit even longer maturities beyond 3 and 4 months, we propose to use a time-dependent parametrization of H in (2.2) in the form of

$$H(t) = H_0 e^{-\kappa t} + H_\infty (1 - e^{-\kappa t}), \quad (5.3)$$

with $H_0, H_\infty, \kappa > 0$ to be calibrated. With this formulation, X remains a Gaussian Ornstein-Uhlenbeck process with time dependent parameters and can also be simulated exactly. The formula for VIX_T^2 remains polynomial in X_T similar to (3.4).

Using time dependent parametrization of H , together with minor tweaks to the stripped forward variance curve using (5.1) and letting the mean reversion speed ε free, we can jointly fit the SPX and VIX surface beyond 1 year, with up to 8 slices for SPX and 6 slices for VIX as illustrated on Figure 5. The calibrated parameters are $\rho = -0.7466$, $(\alpha_0, \alpha_1, \alpha_3, \alpha_5) = (0, 0.0266, 0.2513, 0.00006)$, $H_0 = 0.3176$, $H_\infty = -1.3665$, $\kappa = 1.2$, $\varepsilon = 0.1359$. Figure 6 shows the forward variance curve $\xi_0(t)$ on 23 October 2017 stripped from the market vs. slightly adjusted forward variance as part of the joint calibration, and Figure 7 shows the value of the calibrated H in (5.3) as function of time.

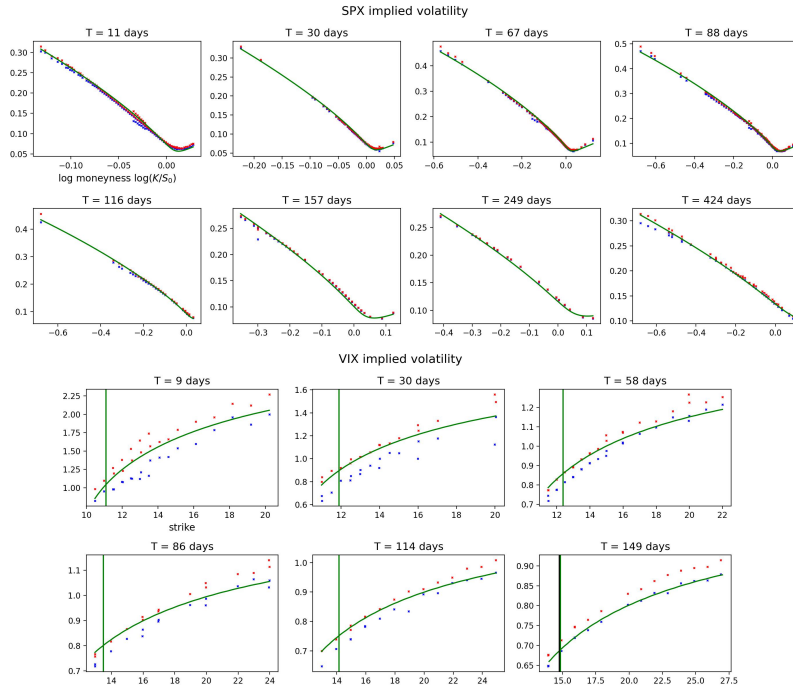


Figure 5: SPX–VIX smiles (bid/ask in blue/red) and VIX futures (vertical black lines) jointly calibrated with our model for time dependent H (full green lines) for 23 October 2017.




Figure 6: The blue line represent the forward variance curve stripped from market data using Carr-Madan log contract formula as in (5.1), the dotted green line is the adjusted forward variance curve as part of the calibration to jointly fit SPX/VIX smiles on 23 October 2017, with the round points representing cubic spline nodes.




Figure 7: Value of H as a function of time as part of the calibration to jointly fit SPX/VIX smiles on 23 October 2017.

Recommendation on the choice of the forward curve ξ_0 and the parameters ε and H for practical use of the model:

- parametric ξ_0 as in (5.2), fixed $\varepsilon = 1/52$ and constant coefficient H in (2.2) for fits of single slice of VIX and two slices of SPX,
- tweaked stripped forward curve ξ_0 , fixed $\varepsilon = 1/52$ and constant H in (2.2) for joint fits on several maturities up to 3 to 4 months,
- tweaked stripped forward curve ξ_0 , letting ε free and time-dependent H in (2.2) as in (5.3) for joint fits on several maturities up to 18 months.

References

[1] Eduardo Abi Jaber and Nathan De Carvalho. Reconciling rough volatility with jumps. *Working paper*, 2022.

[2] Eduardo Abi Jaber, Camille Illand, and Shaun Xiaoyuan Li. Joint SPX–VIX calibration with gaussian polynomial volatility models: deep pricing with quantization hints. *Available at SSRN 4292544*, 2022.

[3] Jan Baldeaux and Alexander Badran. Consistent modelling of vix and equity derivatives using a 3/2 plus jumps model. *Applied Mathematical Finance*, 21(4):299–312, 2014.

[4] Alessandro Bondi, Sergio Pulido, and Simone Scotti. The rough hawkes heston stochastic volatility model. *arXiv preprint arXiv:2210.12393*, 2022.

[5] Peter Carr and Dilip Madan. Towards a theory of volatility trading. *Option Pricing, Interest Rates and Risk Management, Handbooks in Mathematical Finance*, 22(7):458–476, 2001.

- [6] Rama Cont and Thomas Kokholm. A consistent pricing model for index options and volatility derivatives. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 23(2):248–274, 2013.
- [7] J-P Fouque and Yuri F Saporito. Heston stochastic vol-of-vol model for joint calibration of vix and s&p 500 options. *Quantitative Finance*, 18(6):1003–1016, 2018.
- [8] Jean-Pierre Fouque, George Papanicolaou, and K Ronnie Sircar. *Derivatives in financial markets with stochastic volatility*. Cambridge University Press, 2000.
- [9] Jean-Pierre Fouque, George Papanicolaou, Ronnie Sircar, and Knut Solna. Multiscale stochastic volatility asymptotics. *Multiscale Modeling & Simulation*, 2(1):22–42, 2003.
- [10] Jim Gatheral, Paul Jusselin, and Mathieu Rosenbaum. The quadratic rough heston model and the joint s&p 500/vix smile calibration problem. *arXiv preprint arXiv:2001.01789*, 2020.
- [11] Stéphane Goutte, Amine Ismail, and Huyên Pham. Regime-switching stochastic volatility model: estimation and calibration to vix options. *Applied Mathematical Finance*, 24(1):38–75, 2017.
- [12] Lech A Grzelak. On randomization of affine diffusion processes with application to pricing of options on vix and s&p 500. *arXiv preprint arXiv:2208.12518*, 2022.
- [13] Ivan Guo, Gregoire Loeper, Jan Obloj, and Shiyi Wang. Joint modeling and calibration of spx and vix by optimal transport. *SIAM Journal on Financial Mathematics*, 13(1):1–31, 2022.
- [14] Julien Guyon. The joint s&p 500/vix smile calibration puzzle solved. *Risk*, April, 2020.
- [15] Julien Guyon. Dispersion-constrained martingale schrödinger bridges: Joint entropic calibration of stochastic volatility models to s&p 500 and vix smiles. *Available at SSRN 4165057*, 2022.
- [16] Julien Guyon and Jordan Lekeufack. Volatility is (mostly) path-dependent. *Volatility Is (Mostly) Path-Dependent (July 27, 2022)*, 2022.
- [17] Thomas Kokholm and Martin Stisen. Joint pricing of vix and spx options with stochastic volatility and jump models. *The Journal of Risk Finance*, 2015.
- [18] Ryan McCrickerd and Mikko S Pakkanen. Turbocharging monte carlo pricing for the rough bergomi model. *Quantitative Finance*, 18(11):1877–1886, 2018.
- [19] Serguei Mechkov. Fast-reversion limit of the heston model. *Available at SSRN 2418631*, 2015.
- [20] Claudio Pacati, Gabriele Pompa, and Roberto Renò. Smiling twice: the heston++ model. *Journal of Banking & Finance*, 96:185–206, 2018.
- [21] Andrew Papanicolaou and Ronnie Sircar. A regime-switching heston model for vix and s&p 500 implied volatilities. *Quantitative Finance*, 14(10):1811–1827, 2014.
- [22] Sigurd Emil Rømer. Empirical analysis of rough and classical stochastic volatility models to the spx and vix markets. *Quantitative Finance*, 22(10):1805–1838, 2022.
- [23] Mathieu Rosenbaum and Jianfei Zhang. Deep calibration of the quadratic rough heston model. *arXiv preprint arXiv:2107.01611*, 2021.