Pricing volatility derivatives under rough volatility

Blanka Horvath[1] Antoine Jacquier[1] Peter Tankov[2]

[1]Imperial College, London

[2]ENSAE ParisTech, France

Advances in Stochastic Analysis for Risk Modeling

Lumini, November 13-17, 2017

Outline

1. Introduction
2. Pricing variance options: toy example
3. Rough Volterra stochastic volatility
4. Pricing and hedging VIX options
5. Volatility modulated Volterra processes

Volatility is rough!

- Recent studies (Gatheral et al. '14; Bennedsen et al. '16): fractional Brownian motion with $H < \frac{1}{2}$ provides a very good fit to log-volatility time series across many markets and asset classes.
- Since FBM with $H < \frac{1}{2}$ has less regular paths than standard BM, these models have been termed rough volatility models.

Volatility is rough!

- Recent studies (Gatheral et al. '14; Bennedsen et al. '16): fractional Brownian motion with $H < \frac{1}{2}$ provides a very good fit to log-volatility time series across many markets and asset classes.
- Since FBM with $H < \frac{1}{2}$ has less regular paths than standard BM, these models have been termed rough volatility models.
- Rough volatility models perform well for short-term volatility forecasting and describe well the short-term implied volatility smile behavior.

Volatility is rough!

- Recent studies (Gatheral et al. '14; Bennedsen et al. '16): fractional Brownian motion with $H < \frac{1}{2}$ provides a very good fit to log-volatility time series across many markets and asset classes.
- Since FBM with $H < \frac{1}{2}$ has less regular paths than standard BM, these models have been termed rough volatility models.
- Rough volatility models perform well for short-term volatility forecasting and describe well the short-term implied volatility smile behavior.
- FBM with $H \neq \frac{1}{2}$ is not a Markov process nor a semimartingale.
- This talk: efficient pricing and hedging algorithms for volatility/variance options in such models; calibration to VIX smiles.

ATM skew under rough volatility

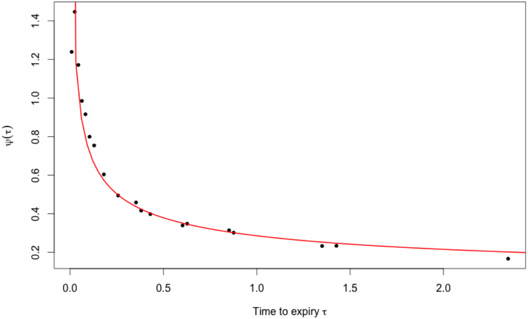

ATM skew $\frac{\partial C}{\partial K}\Big|_{K=S}$ of S&P options on Aug 14, 2013, and fit by $\psi(\tau) = A\tau^{-0.407}$.
This behavior is compatible with FBM volatility with $H = 0.093$.

Source: Bayer, Friz, Gatheral, Pricing under rough volatility (2015).

Volatility forecasting

Forecasting formula (Nutzman and Poor, 00):

$$\mathbb{E}[W^H_{t+\theta}|\mathcal{F}_t] = W^H_t + \frac{\cos(\pi H)}{\pi}\theta^{H+1/2}\int_0^\infty \frac{W^H_{t-s} - W^H_t}{s^{H+1/2}(s+\theta)}ds.$$

Forecasting performance is horizon-independent:

$$\frac{\mathbb{E}[(W^H_{t+\theta} - \mathbb{E}[W^H_{t+\theta}|\mathcal{F}_t])^2]}{\mathbb{E}[(W^H_{t+\theta} - W^H_t)^2]} = \frac{\Gamma\left(\frac{3}{2} - H\right)}{\Gamma(2-2H)\Gamma\left(H + \frac{1}{2}\right)}$$

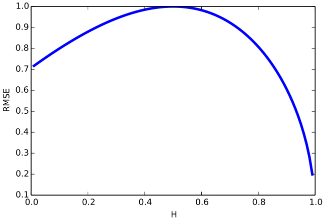

Outline

Toy example

Assume the instantaneous volatility is

$$\sigma_t = \sigma e^{X_t},$$

where X is general centered Gaussian under $\mathbb{Q}$; let $\mathcal{F}_s^0 = \sigma(X_r, r \leq s)$ and $\mathcal{F}_s = \cup_{s>t} \mathcal{F}_s^0$.

Toy example

Assume the instantaneous volatility is

$$\sigma_t = \sigma e^{X_t},$$

where X is general centered Gaussian under $\mathbb{Q}$; let $\mathcal{F}^0_s = \sigma(X_r, r \leq s)$ and $\mathcal{F}_s = \cup_{s>t}\mathcal{F}^0_s$.

Fix a time horizon T, and let $Z_t = \mathbb{E}[X_T|\mathcal{F}_t]$. Then Z_t is a Gaussian martingale and thus a PII characterized by

$$c(t) = \mathbb{E}[Z_t^2] = \mathbb{E}[\mathbb{E}[X_T|\mathcal{F}_t]^2].$$

Assume that $c(t)$ is continuous so that Z_t is continuous.

Toy example

Assume the instantaneous volatility is

$$\sigma_t = \sigma e^{X_t},$$

where X is general centered Gaussian under $\mathbb{Q}$; let $\mathcal{F}_s^0 = \sigma(X_r, r \leq s)$ and $\mathcal{F}_s = \cup_{s>t} \mathcal{F}_s^0$.

Fix a time horizon T, and let $Z_t = \mathbb{E}[X_T|\mathcal{F}_t]$. Then Z_t is a Gaussian martingale and thus a PII characterized by

$$c(t) = \mathbb{E}[Z_t^2] = \mathbb{E}[\mathbb{E}[X_T|\mathcal{F}_t]^2].$$

Assume that $c(t)$ is continuous so that Z_t is continuous.

Introduce the forward variance

$$\xi_t = \mathbb{E}[\sigma_T^2|\mathcal{F}_t] = \sigma^2 \mathbb{E}[e^{2X_T}|\mathcal{F}_t] = Ce^{2(Z_t - c(t))}.$$

Toy example

The time-t price of a call option on instantaneous forward variance

$$\mathbb{E}[(\xi_{T_0} - K)^+ | \mathcal{F}_t].$$

Since $(\xi_t)_{t \geq 0}$ is a continuous log-normal martingale, we can write,

$$P_t = \mathbb{E}[(\xi_{T_0} - K)^+ | \mathcal{F}_t] = P(t, \xi_t),$$

where P is a deterministic function given by

$$P(t, x) = \mathbb{E}[(xe^{Z - \frac{1}{2}\mathrm{Var} Z} - K)^+]$$

and Z is a centered Gaussian random variable with variance $4(c(T_0) - c(t))$.

Toy example

By the Black's formula,

$$P_t = \mathbb{E}[(\xi_{T_0} - K)^+ | \mathcal{F}_t] = \xi_t N(d_t^1) - KN(d_t^2),$$

where N is the standard normal distribution function and

$$d_t^{1,2} = \frac{\log \frac{\xi_t}{K} \pm 4(c(T_0) - c(t))}{2\sqrt{c(T_0) - c(t)}}.$$

Toy example

By the Black's formula,

$$P_t = \mathbb{E}[(\xi_{T_0} - K)^+ | \mathcal{F}_t] = \xi_t N(d_t^1) - KN(d_t^2),$$

where N is the standard normal distribution function and

$$d_t^{1,2} = \frac{\log \frac{\xi_t}{K} \pm 4(c(T_0) - c(t))}{2\sqrt{c(T_0) - c(t)}}.$$

Applying the Itô formula, we get,

$$dP_t = N(d_t^1) d\xi_t.$$

- The forward variance option may be hedged perfectly by a portfolio containing the "instantaneous" variance swap and the risk-free asset

Outline

Gaussian Volterra processes

- The FBM is self-similar $\Rightarrow$ it exhibits the same behavior at different time scales, which is not always the case in the data.
- The model is supposed to describe the whole forward variance curve $\Rightarrow$ one stochastic factor is not enough

Gaussian Volterra processes

- The FBM is self-similar $\Rightarrow$ it exhibits the same behavior at different time scales, which is not always the case in the data.
- The model is supposed to describe the whole forward variance curve $\Rightarrow$ one stochastic factor is not enough
- Gaussian Volterra processes (Bennedsen et al. '16) allow for a general kernel:

$$\sigma_t = \sigma e^{X_t} \quad \text{with} \quad X_t = \int_{-\infty}^{t} g(t,s)^\top dW_s,$$

where W is a d-dimensional BM on $\mathbb{R}$ with respect to $\mathbb{F} \equiv (\mathcal{F}_t)_{t\in\mathbb{R}}$, and the function g is such that

$$\int_{-\infty}^{t} \|g(t,s)\|^2 ds < \infty, \quad \forall t \geq 0.$$

Rough Volterra stochastic volatility

The forward variance $\xi_t(u) = \mathbb{E}[\sigma_u^2|\mathcal{F}_t]$ has explicit martingale dynamics

$$d\xi_t(u) = 2\xi_t(u)g(u,t)^\top dW_t.$$

Rough Volterra stochastic volatility

The forward variance $\xi_t(u) = \mathbb{E}[\sigma_u^2|\mathcal{F}_t]$ has explicit martingale dynamics

$$d\xi_t(u) = 2\xi_t(u)g(u,t)^\top dW_t.$$

Example: "rough Bergomi" model: $\sigma_t = \sigma e^{\alpha W_t^H}$, where W^H is a FBM with Hurst parameter H (see Bayer, Friz and Gatheral '15). This model corresponds to

$$g(t,s) = \alpha \mathbf{1}_{s<0}[(t-s)^{H-\frac{1}{2}} - (-s)^{H-\frac{1}{2}}] + \alpha \mathbf{1}_{s\geq 0}(t-s)^{H-\frac{1}{2}}.$$

The dynamics of the forward variance $\xi_t(u)$ is

$$\frac{d\xi_t(T)}{\xi_t(T)} = 2\alpha(T-t)^{H-\frac{1}{2}} dW_t, \quad 0 \leq t \leq T$$

Outline

VIX options

- The VIX index, quoted by CBOE since 2001 is the index of implied volatility of S&P 500 options, computed using the variance swap formula
- Options and futures on the VIX are liquidly quoted
- In our model the VIX index at time T takes the form

$$\sqrt{-\frac{2}{\Theta}\mathbb{E}\left[\log\frac{S_{T+\Theta}}{S_T}\middle|\mathcal{F}_T\right]} = \sqrt{\frac{1}{\Theta}\int_T^{T+\Theta}\xi_T(u)du}$$

where $\Theta = 1$ month.

- We consider a VIX option with pay-off at time T given by

$$f\left(\frac{1}{\Theta}\int_T^{T+\Theta}\xi_T(u)du\right).$$

VIX options

The time-t price of a VIX option is given by

$$P_t = \mathbb{E}\left[f\left(\frac{1}{\Theta}\int_T^{T+\Theta}\xi_T(u)du\right)\Big|\mathcal{F}_t\right] = F(t, \xi_t(u)_{T\leq u\leq T+\Theta}),$$

where F is a deterministic from $[0, T]\times H$ with $H = L^2([T, T+\Theta])$ to $\mathbb{R}$:

$$F(t,x) = \mathbb{E}\left[f\left(\frac{1}{\Theta}\int_T^{T+\Theta}x(u)\mathcal{E}_{t,T}(u)du\right)\right],$$

where

$$\mathcal{E}_{t,T}(u) := \mathcal{E}\left(2\int_t^\cdot g(u,s)^\top dW_s\right)_T.$$

Hedging VIX options

- Related ongoing work by Masaaki Fukasawa (presented in NY on Oct 14)
- Hedging options in stock price models modulated by FBM was also discussed by Djehiche & Eddahbi (2001) using Malliavin calculus methods

Theorem

Let the function f be differentiable with f' piecewise continuous and bounded. Then the option price P_t admits the martingale representation

$$P_T = P_t + 2\int_t^T \int_T^{T+\Theta} D_x F(s,\xi_s)(u)\xi_s(u)g(u,s)^\top dW_s,$$

where the Frechet derivative $D_x F$ is given explicitly by

$$D_x F(t,x)(v) = \mathbb{E}\left[f'\left(\int_T^{T+\Theta} x(u)\mathcal{E}_{t,T}(u)du\right)\mathcal{E}_{t,T}(v)\right]$$

Pricing VIX options by Monte Carlo

- To price a VIX option, only need to discretize the integral $\int_T^{T+\Theta} \xi_T(u)du$
- We consider two different discretization schemes: rectangle scheme

$$F_n(t,x) = \mathbb{E}\left[f\left(\frac{1}{\Theta}\sum_{i=0}^{n-1}\xi_i^n \mathcal{E}_{t,T}(t_i^n)\right)\right],$$

where $\xi_i^n = \int_{t_i^n}^{t_{i+1}^n} x(u)du$, and trapezoidal scheme

$$\widehat{F}_n(t,x) = \mathbb{E}\left[f\left(\frac{1}{\Theta}\sum_{i=0}^{n-1}\int_{t_i^n}^{t_{i+1}^n} x(u)(\theta^n(u)\mathcal{E}_{t,T}(t_i^n) + (1-\theta^n(u))\mathcal{E}_{t,T}(t_{i+1}^n))du\right)\right]$$

where $\theta^n(u) = \frac{t_{i+1}^n - u}{t_{i+1}^n - t_i^n}$.

Convergence rates

- Convergence rates depend on the singularity of the kernel function $g(u,t)$ when $u \to t$.
- For the rough Bergomi model with Hurst parameter H, for the rectangle scheme with $t_i^n = T + \Theta\frac{i}{n}$,

$$|F(t,x) - F_n(t,x)| \leq \frac{C}{n}$$

Convergence rates

- Convergence rates depend on the singularity of the kernel function $g(u,t)$ when $u \to t$.
- For the rough Bergomi model with Hurst parameter H, for the rectangle scheme with $t_i^n = T + \Theta \frac{i}{n}$,

$$|F(t,x) - F_n(t,x)| \leq \frac{C}{n}$$

- For the trapezoidal scheme with $t_i^n = T + \Theta \left(\frac{i}{n}\right)^\eta$, $\eta(H+1) > 2$,

$$|F(t,x) - \widehat{F}_n(t,x)| \leq \frac{C}{n^2}$$

Convergence of the Monte Carlo estimator

Payoff $f(x) = (\sqrt{x} - K)^+$ with $K = 0.2$ (Lipschitz)

Parameters: $\alpha = 0.2$, $H = 0.1$, flat forward variance with $\xi = 0.2$, $t = 0$, $T = 1.0$, $\Theta = .1$, 50000 MC runs

Slopes of log-log graph: 1.15 (1st order) and 2.76

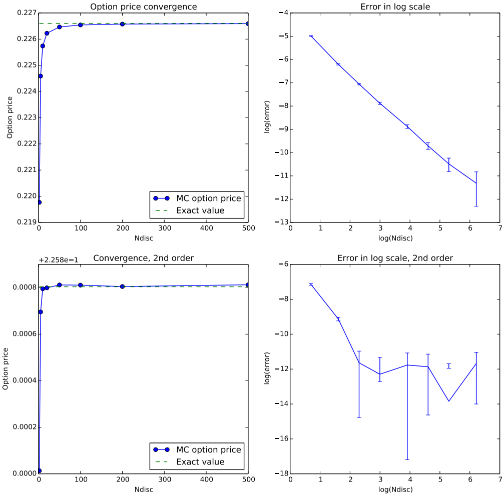

VIX Implied volatility smiles in the rough Bergomi model

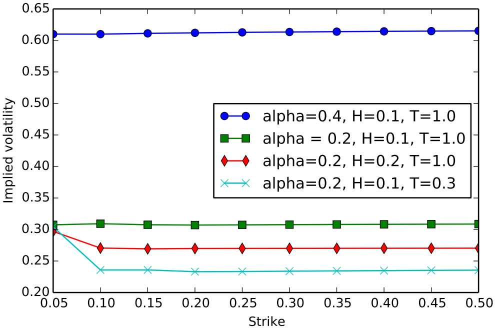

The implied volatility of VIX options is defined assuming that VIX future is log-normal and using the model-generated VIX future as initial value.

Log-normal model: very small correction with respect to constant volatility price $\Rightarrow$ inconsistency with market VIX smiles (Bayer, Friz and Gatheral '15).

VIX implied volatility smiles: the market data

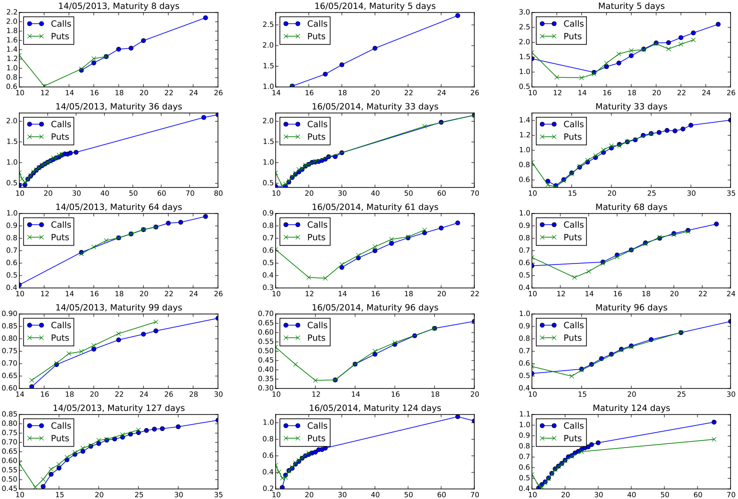

Outline

Enhancing rough volatility models with VIX smiles

- In Gaussian Volterra model the forward volatility is log-normal
  $\Rightarrow$ flat smile of VIX options

Enhancing rough volatility models with VIX smiles

- In Gaussian Volterra model the forward volatility is log-normal
  $\Rightarrow$ flat smile of VIX options
- In an ongoing work (Bachelier seminar, November 10, 2017), Stefano De Marco presents an approach inspired by the work of Bergomi (2008):

$$\xi_t(T) = \xi_0(T) f^T(t, x_t^T),$$

where x_t^T is the Gaussian Volterra process, and f^T is a smooth function to be calibrated to VIX smiles

Enhancing rough volatility models with VIX smiles

- In Gaussian Volterra model the forward volatility is log-normal
  $\Rightarrow$ flat smile of VIX options
- In an ongoing work (Bachelier seminar, November 10, 2017), Stefano De Marco presents an approach inspired by the work of Bergomi (2008):

$$\xi_t(T) = \xi_0(T) f^T(t, x_t^T),$$

  where x_t^T is the Gaussian Volterra process, and f^T is a smooth function to be calibrated to VIX smiles
- We follow an alternative route, based on modulated Volterra processes (Barndorff-Nielsen, Benth and Veraart '12)

Volatility modulated Gaussian Volterra processes

To allow for VIX smiles, introduce volatility modulated Volterra processes following Barndorff-Nielsen, Benth and Veraart (2012):

$$\sigma_t = e^{X_t}, \qquad X_t = \int_{-\infty}^{t} \sqrt{\Gamma_s} g(t,s)^\top dW_s,$$

where Γ is a positive affine process independent of W

- Can introduce different modulators for different factors
- Γ can be a CIR process, or a positive Lévy-driven Ornstein-Uhlenbeck process

Volatility modulated Gaussian Volterra processes

Once again, switch to forward variance $\xi_t(u) = \mathbb{E}[\sigma_u^2|\mathcal{F}_t]$

$$\xi_t(u) = \exp\left(2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s) dW_s\right) \mathbb{E}\left[\exp\left(2\int_t^u \|g(u,s)\|^2 \Gamma_s ds\right) \Big| \mathcal{F}_t\right]$$

$$= \exp\left(2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s) dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right)$$

where ψ and ψ_0 are known coefficients coming from the affine structure of Γ

Volatility modulated Gaussian Volterra processes

Once again, switch to forward variance $\xi_t(u) = \mathbb{E}[\sigma_u^2|\mathcal{F}_t]$

$$\xi_t(u) = \exp\left(2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s\right)\mathbb{E}\left[\exp\left(2\int_t^u \|g(u,s)\|^2\Gamma_s ds\right)\Big|\mathcal{F}_t\right]$$

$$= \exp\left(2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right)$$

where ψ and ψ_0 are known coefficients coming from the affine structure of Γ

Explicit Markov martingale dynamics for forward variance curve:

$$\frac{d\xi_t(u)}{\xi_{t-}(u)} = 2\sqrt{\Gamma_t} g(u,t)^\top dW_t + \psi_0(t,u)d\Gamma_t^c + \int_{\mathbb{R}} (e^{\psi_0(t,u)z} - 1)\widetilde{J}_\Gamma(dt \times dz),$$

where Γ^c is the diffusion part of Γ and $\widetilde{J}_\Gamma$ its compensated jump measure

$\Rightarrow$ can have a highly skewed distribution even if Γ is independent from W

Example

Let Γ be Lévy-driven positive Ornstein-Uhlenbeck process:

$$d\Gamma_t = -\lambda \Gamma_t dt + dL_t, \qquad \mathbb{E}[e^{uL_t}] = e^{t\psi(u)}$$

Then,

$$\mathbb{E}\left[\exp\left(2\int_t^u \|g(u,s)\|^2 \Gamma_s ds\right) \Big| \mathcal{F}_t\right] = e^{\Gamma_t \psi_0(t,u) + \psi(t,u)}$$

with

$$\psi_0(t,u) = 2\int_t^u e^{-\lambda(s-t)} \|g(u,s)\|^2 ds$$

$$\psi(t,u) = \int_t^u \psi\left(2\int_s^t \|g(u,r)\|^2 e^{-\lambda(r-s)} dr\right) ds$$

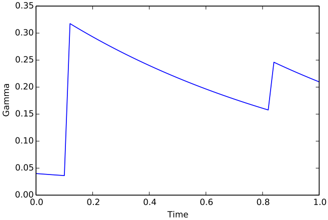

Pricing VIX options: approximation

For pricing a VIX option, we need to simulate

$$\frac{1}{\Theta}\int_T^{T+\Theta} \xi_t(u)du$$
$$= \frac{1}{\Theta}\int_T^{T+\Theta} du \exp\left(2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right)$$

Pricing VIX options: approximation

For pricing a VIX option, we need to simulate

$$\frac{1}{\Theta}\int_T^{T+\Theta} \xi_t(u)du$$

$$= \frac{1}{\Theta}\int_T^{T+\Theta} du \exp\left(2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right)$$

This can be approximated by

$$\exp\left(\frac{1}{\Theta}\int_T^{T+\Theta} du \left\{2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right\}\right)$$

Pricing VIX options: approximation

For pricing a VIX option, we need to simulate

$$\frac{1}{\Theta}\int_T^{T+\Theta} \xi_t(u)du$$
$$= \frac{1}{\Theta}\int_T^{T+\Theta} du \exp\left(2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right)$$

This can be approximated by

$$\exp\left(\frac{1}{\Theta}\int_T^{T+\Theta} du \left\{2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right\}\right)$$

The characteristic function of

$$\frac{1}{\Theta}\int_T^{T+\Theta} du \left\{2\int_{-\infty}^t \sqrt{\Gamma_s} g(u,s)dW_s + \psi_0(t,u)\Gamma_t + \psi(t,u)\right\}$$

is known, leading to explicit formulas for approximate pricing / calibration

Pricing VIX options: Monte Carlo

- The forward variance curve is conditionnally log-normal given the initial curve and the trajectory of Γ_t.
- Discretization is only necessary for the period over which the variance is integrated (1 month for VIX options) but not over the lifespan of the option.
- We fix the discretization grid $T = t_0^n < t_1^n < \cdots < t_n^n = T + \Theta$ and simulate Γ_T and the covariances

$$C_{ij} = \int_t^T \Gamma_s g(t_i^n, s)^\top g(t_j^n, s) ds$$

Pricing VIX options: Monte Carlo

- The forward variance curve is conditionnally log-normal given the initial curve and the trajectory of Γ_t.
- Discretization is only necessary for the period over which the variance is integrated (1 month for VIX options) but not over the lifespan of the option.
- We fix the discretization grid $T = t_0^n < t_1^n < \dots < t_n^n = T + \Theta$ and simulate Γ_T and the covariances

$$C_{ij} = \int_t^T \Gamma_s g(t_i^n, s)^\top g(t_j^n, s) ds$$

- In the case when Γ is a Lévy-driven OU process with finite jump intensity, this simulation is done without error.
- In the last step, $\mathcal{E}_{t,T}(t_i^n)$ are simulated and the option pay-off is computed.

Calibration of the modulated rough Bergomi model

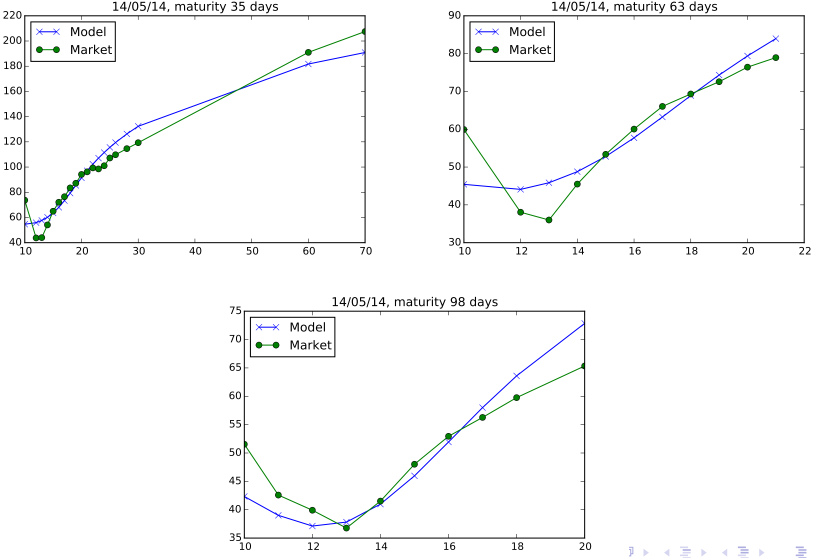

Conclusions

- Rough volatility models reproduce both the observed dynamics of volatility and the short term implied volatility skew
- Working with forward variance allows to recover martingale techniques, develop efficient Monte Carlo and asymptotic pricing and hedging methods
- However, log-normal rough volatility models remain inconsistent with VIX option smiles
- A volatility modulated model allows to generate VIX smiles consistent with the market while preserving some simplicity of log-normal modeling