

Rough volatility: An overview

Jim Gatheral



Financial Engineering Practitioners Seminar,
Columbia University,
Monday January 22, 2018

Term structure of SPX ATM skew as of 15-Sep-2005




Figure 2: Term structure of ATM skew as of 15-Sep-2005, with power law fit $\tau^{-0.44}$ superimposed in red.

Stylized facts

- Although the levels and orientations of the volatility surfaces change over time, their rough shape stays very much the same.
 - It's then natural to look for a time-homogeneous model.
- The term structure of ATM volatility skew

$$\psi(\tau) \sim \frac{1}{\tau^\alpha}$$

with $\alpha \in (0.3, 0.5)$.

Motivation for Rough Volatility I: Better fitting stochastic volatility models

- Conventional stochastic volatility models generate volatility surfaces that are inconsistent with the observed volatility surface.
 - In stochastic volatility models, the ATM volatility skew is constant for short dates and inversely proportional to T for long dates.
 - Empirically, we find that the term structure of ATM skew is proportional to $1/T^\alpha$ for some $0 < \alpha < 1/2$ over a very wide range of expirations.
- The conventional solution is to introduce more volatility factors, as for example in the DMR and Bergomi models.
 - One could imagine the power-law decay of ATM skew to be the result of adding (or averaging) many sub-processes, each of which is characteristic of a trading style with a particular time horizon.

Bergomi Guyon

- Define the forward variance curve $\xi_t(u) = \mathbb{E}[v_u | \mathcal{F}_t]$.
- According to [BG12], in the context of a variance curve model, implied volatility may be expanded as

$$\sigma_{\text{BS}}(k, T) = \sigma_0(T) + \sqrt{\frac{w}{T}} \frac{1}{2w^2} C^{\times\xi} k + O(\eta^2) \quad (1)$$

where η is volatility of volatility, $w = \int_0^T \xi_0(s) ds$ is total variance to expiration T , and

$$C^{\times\xi} = \int_0^T dt \int_t^T du \frac{\mathbb{E}[dx_t d\xi_t(u)]}{dt}. \quad (2)$$

- Thus, given a stochastic volatility model written in forward variance form, we can easily (at least in principle) compute this smile approximation.

The Bergomi model

- The n -factor Bergomi variance curve model reads:

$$\xi_t(u) = \xi_0(u) \exp \left\{ \sum_{i=1}^n \eta_i \int_0^t e^{-\kappa_i(t-s)} dW_s^{(i)} + \text{drift} \right\}. \quad (3)$$

- The Bergomi model generates a term structure of volatility skew $\psi(\tau)$ that is something like

$$\psi(\tau) = \sum_i \frac{1}{\kappa_i \tau} \left\{ 1 - \frac{1 - e^{-\kappa_i \tau}}{\kappa_i \tau} \right\}.$$

- This functional form is related to the term structure of the autocorrelation function.
- Which is in turn driven by the exponential kernel in the exponent in (3).

Tinkering with the Bergomi model

- Empirically, $\psi(\tau) \sim \tau^{-\alpha}$ for some α .
- It's tempting to replace the exponential kernels in (3) with a power-law kernel.
- This would give a model of the form

$$\xi_t(u) = \xi_0(u) \exp \left\{ \eta \int_0^t \frac{dW_s}{(t-s)^\gamma} + \text{drift} \right\}$$

which looks similar to

$$\xi_t(u) = \xi_0(u) \exp \left\{ \eta W_t^H + \text{drift} \right\}$$

where W_t^H is fractional Brownian motion.

History of fractional stochastic volatility models

More formally, the model

$$\xi_t(u) = \xi_0(u) \exp \left\{ \eta \int_0^t \frac{dW_s}{(t-s)^\gamma} + \text{drift} \right\}$$

belongs to a larger class of fractional stochastic volatility models that was originally shown by Alòs et al. in [ALV07] and then by Fukasawa in [Fuk11] to generate a short-dated ATM skew of the form

$$\psi(\tau) \sim \frac{1}{\tau^\gamma}$$

with $\gamma = \frac{1}{2} - H$ and $0 < H < \frac{1}{2}$.

Motivation for Rough Volatility II: Power-law scaling of the historical volatility process

- The Oxford-Man Institute of Quantitative Finance makes historical realized variance (RV) estimates freely available at <http://realized.oxford-man.ox.ac.uk>. These estimates are updated daily.
- Using daily RV estimates as proxies for instantaneous variance, we may investigate the time series properties of v_t empirically.

SPX realized variance from 2000 to 2016

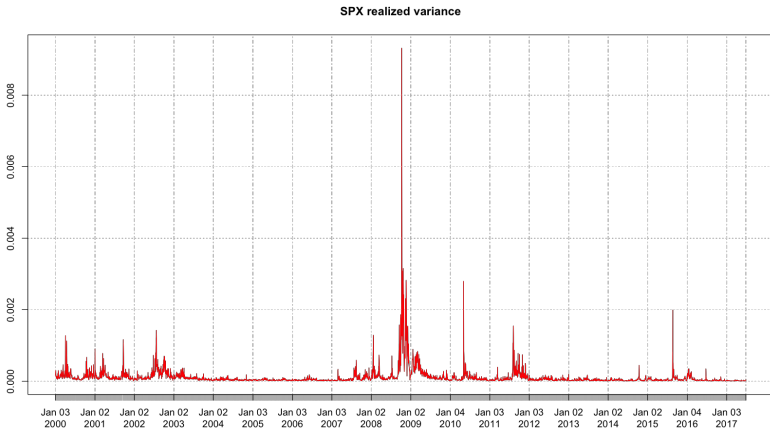


Figure 3: KRV estimates of SPX realized variance from 2000 to 2017.

The smoothness of the volatility process

- For $q \geq 0$, we define the q th sample moment of differences of log-volatility at a given lag Δ ¹:

$$m(q, \Delta) = \langle |\log \sigma_{t+\Delta} - \log \sigma_t|^q \rangle$$

- For example

$$m(2, \Delta) = \langle (\log \sigma_{t+\Delta} - \log \sigma_t)^2 \rangle$$

is just the sample variance of differences in log-volatility at the lag Δ .

¹ $\langle \cdot \rangle$ denotes the sample average.

Scaling of $m(q, \Delta)$ with $\log \Delta$




Figure 4: $\log m(q, \Delta)$ as a function of $\log \Delta$, SPX.

Scaling of ζ_q with q




Figure 5: Scaling of ζ_q with q .

Monofractal scaling result

- From the log-log plot Figure 4, we see that for each q , $m(q, \Delta) \propto \Delta^{\zeta_q}$.
- And from Figure 5 the monofractal scaling relationship

$$\zeta_q = qH$$

with $H \approx 0.13$.

- Note also that our estimate of H is biased high because we proxied instantaneous variance v_t with its average over each day $\frac{1}{T} \int_0^T v_t dt$, where T is one day.
- On the other hand, the time series of realized variance is noisy and this causes our estimate of H to be biased low.
- A time series of H for SPX following the methodology of [BLP16] is shown in the next figure.

The time series of $\hat{\alpha} = H - \frac{1}{2}$ for SPX




Figure 10: Half year rolling-window estimates of α on the realized variance measures of the daily volatility by variogram OLS regression (3.10) with $m = 3$. The pink area is the 95% confidence interval by bootstrap method with $B = 999$. The four vertical dashed blue lines indicate four periods of market turmoil: Lehman Brothers filing for bankruptcy, the Flash Crash, the first bailout during Greek debt crisis and the Brexit referendum.

Distributions of $(\log \sigma_{t+\Delta} - \log \sigma_t)$ for various lags Δ

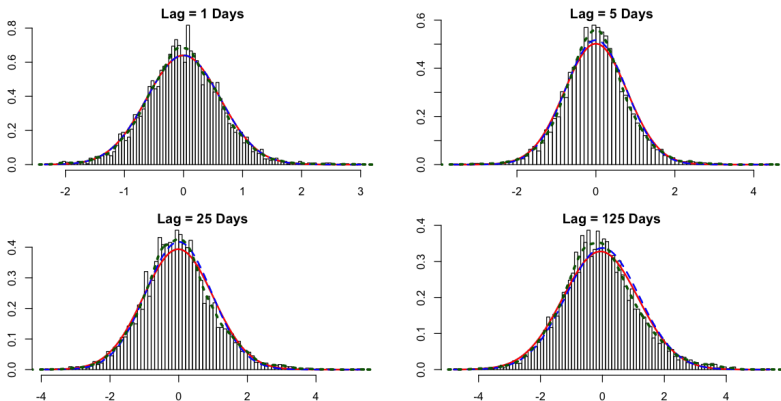


Figure 6: Histograms of $(\log \sigma_{t+\Delta} - \log \sigma_t)$ for various lags Δ ; normal fit in red; $\Delta = 1$ normal fit scaled by $\Delta^{0.14}$ in blue.

Estimated H for all indices

Estimating the relationship

$$\langle (\log \sigma_{t+\Delta} - \log \sigma_t)^2 \rangle = \nu^2 \Delta^{2H}$$

for all 21 indices in the Oxford-Man dataset yields:

| Index | H | ν |
|-------------|------|-------|
| SPX2.rk | 0.13 | 0.32 |
| FTSE2.rk | 0.14 | 0.27 |
| N2252.rk | 0.11 | 0.33 |
| GDAXI2.rk | 0.15 | 0.28 |
| RUT2.rk | 0.12 | 0.33 |
| AORD2.rk | 0.08 | 0.36 |
| DJI2.rk | 0.13 | 0.32 |
| IXIC2.rk | 0.13 | 0.30 |
| FCHI2.rk | 0.13 | 0.29 |
| HSI2.rk | 0.10 | 0.28 |
| KS11.rk | 0.12 | 0.28 |
| AEX.rk | 0.14 | 0.29 |
| SSMI.rk | 0.18 | 0.22 |
| IBEX2.rk | 0.13 | 0.28 |
| NSEI.rk | 0.11 | 0.32 |
| MXX.rk | 0.09 | 0.33 |
| BVSP.rk | 0.11 | 0.31 |
| GSPTSE.rk | 0.12 | 0.31 |
| STOXX50E.rk | 0.12 | 0.34 |
| FTSTI.rk | 0.13 | 0.23 |
| FTSEMIB.rk | 0.13 | 0.30 |

Universality?

- In [GJR18], we compute daily realized variance estimates over one hour windows for DAX and Bund futures contracts, finding similar scaling relationships.
- We have also checked that Gold and Crude Oil futures scale similarly.
 - Although the increments $(\log \sigma_{t+\Delta} - \log \sigma_t)$ seem to be fatter tailed than Gaussian.
- In [BLP16] Bennedsen et al., estimate volatility time series for more than five thousand individual US equities, finding rough volatility in every case.

A natural model of realized volatility

- Distributions of differences in the log of realized volatility are close to Gaussian.
 - This motivates us to model σ_t as a lognormal random variable.
- Moreover, the scaling property of variance of RV differences suggests the model:

$$\log \sigma_{t+\Delta} - \log \sigma_t = \nu \left(W_{t+\Delta}^H - W_t^H \right) \quad (4)$$

where W^H is fractional Brownian motion.

- In [GJR18], we refer to a stationary version of (4) as the RFSV (for Rough Fractional Stochastic Volatility) model.

Fractional Brownian motion (fBm)

- *Fractional Brownian motion* (fBm) $\{W_t^H; t \in \mathbb{R}\}$ is the unique Gaussian process with mean zero and autocovariance function

$$\mathbb{E} \left[W_t^H W_s^H \right] = \frac{1}{2} \left\{ |t|^{2H} + |s|^{2H} - |t - s|^{2H} \right\}$$

where $H \in (0, 1)$ is called the *Hurst index* or parameter.

- In particular, when $H = 1/2$, fBm is just Brownian motion.
 - If $H > 1/2$, increments are positively correlated.
 - If $H < 1/2$, increments are negatively correlated.

Apparent fractality of the volatility time series




Figure 7: Volatility of SPX (above) and in the RFSV model (below).

Remarks on the comparison

- The qualitative features of simulated and actual graphs look very similar.
 - Persistent periods of high volatility alternate with low volatility periods.
- $H \sim 0.1$ generates very rough looking sample paths (compared with $H = 1/2$ for Brownian motion).
 - Hence *rough volatility*.
- On closer inspection, we observe fractal-type behavior.
 - The graph of volatility over a small time period looks like the same graph over a much longer time period.
- This feature of volatility has been investigated both empirically and theoretically in, for example, [BM03].
 - In particular, their Multifractal Random Walk (MRW) is related to a limiting case of the RSFV model as $H \rightarrow 0$.

Microstructural foundation for rough volatility

In [JR16], El Euch, Fukasawa and Rosenbaum consider a generalization of a simple model of price dynamics in terms of Hawkes processes due to Bacry et al. ([BM14]) with the following properties:

- Reflecting the high degree of endogeneity of the market, the L^1 norm of the kernel matrix is close to one (nearly unstable).
- No drift in the price process imposes a relationship between buy and sell kernels.
- Liquidity asymmetry: The average impact of a sell order is greater than the impact of a buy order.
- Splitting of metaorders motivates power-law decay of the Hawkes kernels $\varphi(\tau) \sim \tau^{-(1+\alpha)}$ (empirically $\alpha \approx 0.6$).

The scaling limit of the price model

They construct a sequence of such Hawkes processes suitably rescaled in time and space that converges in law to a Rough Heston process of the form

$$\frac{dS_t}{S_t} = \sqrt{v_t} dZ_t$$

$$v_t = v_0 + \frac{\lambda}{\Gamma(\alpha)} \int_0^t \frac{\theta - v_s}{(t-s)^{1-\alpha}} ds + \frac{\lambda \nu}{\Gamma(\alpha)} \int_0^t \frac{\sqrt{v_s} dW_s}{(t-s)^{1-\alpha}}$$

with

$$d\langle Z, W \rangle_t = \rho dt.$$

- The correlation ρ is related to a liquidity asymmetry parameter.
- Rough volatility can thus be understood as relating to the persistence of order flow and the high degree of endogeneity of liquid markets.



The Rough Heston characteristic function

Define the fractional integral and differential operators:

$$I^{1-\alpha}f(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t \frac{f(s)}{(t-s)^\alpha} ds; \quad D^\alpha f(t) = \frac{d}{dt} I^{1-\alpha}f(t).$$

Remarkably, in [ER16], El Euch and Rosenbaum compute the following expression for the characteristic function of the Rough Heston model:

$$\phi_t(u) = \exp \left\{ \theta \lambda \int_0^t h(u, s) ds + v_0 I^{1-\alpha} h(u, t) \right\}$$

where $h(u, \cdot)$ solves the fractional Riccati equation

$$D^\alpha h(u, s) = -\frac{1}{2} u(u+i) + \lambda(i\rho\nu u - 1)h(u, s) + \frac{(\lambda\nu)^2}{2} h^2(u, s).$$

The forecast formula

- Thus, we obtain

Variance forecast formula

$$\mathbb{E}^{\mathbb{P}} [v_{t+\Delta} | \mathcal{F}_t] = \exp \left\{ \mathbb{E}^{\mathbb{P}} [\log(v_{t+\Delta}) | \mathcal{F}_t] + 2 c \nu^2 \Delta^{2H} \right\} \quad (5)$$

where

$$\begin{aligned} & \mathbb{E}^{\mathbb{P}} [\log v_{t+\Delta} | \mathcal{F}_t] \\ &= \frac{\cos(H\pi)}{\pi} \Delta^{H+1/2} \int_{-\infty}^t \frac{\log v_s}{(t-s+\Delta)(t-s)^{H+1/2}} ds. \end{aligned}$$

- [BLP16] confirm that this forecast outperforms the best performing existing alternatives such as HAR, at least at daily or higher timescales.

Pricing under rough volatility

Once again, the data suggests the following model for volatility under the real (or historical or physical) measure \mathbb{P} :

$$\log \sigma_t = \nu W_t^H.$$

Let $\gamma = \frac{1}{2} - H$. We choose the Mandelbrot-Van Ness representation of fractional Brownian motion W^H as follows:

$$W_t^H = C_H \left\{ \int_{-\infty}^t \frac{dW_s^{\mathbb{P}}}{(t-s)^\gamma} - \int_{-\infty}^0 \frac{dW_s^{\mathbb{P}}}{(-s)^\gamma} \right\}$$

where the choice

$$C_H = \sqrt{\frac{2H\Gamma(3/2-H)}{\Gamma(H+1/2)\Gamma(2-2H)}}$$

ensures that

$$\mathbb{E} \left[W_t^H W_s^H \right] = \frac{1}{2} \left\{ t^{2H} + s^{2H} - |t-s|^{2H} \right\}.$$

Pricing under rough volatility

Then

$$\begin{aligned} & \log v_u - \log v_t \\ = & \nu C_H \left\{ \int_t^u \frac{1}{(u-s)^\gamma} dW_s^{\mathbb{P}} + \int_{-\infty}^t \left[\frac{1}{(u-s)^\gamma} - \frac{1}{(t-s)^\gamma} \right] dW_s^{\mathbb{P}} \right\} \\ =: & 2\nu C_H [M_t(u) + Z_t(u)]. \end{aligned} \quad (6)$$

- Note that $\mathbb{E}^{\mathbb{P}} [M_t(u) | \mathcal{F}_t] = 0$ and $Z_t(u)$ is \mathcal{F}_t -measurable.
- To price options, it would seem that we would need to know \mathcal{F}_t , the entire history of the Brownian motion W_s for $s < t$!

Pricing under \mathbb{P}

Let

$$\tilde{W}_t^{\mathbb{P}}(u) := \sqrt{2H} \int_t^u \frac{dW_s^{\mathbb{P}}}{(u-s)^\gamma}$$

With $\eta := 2\nu C_H / \sqrt{2H}$ we have $2\nu C_H M_t(u) = \eta \tilde{W}_t^{\mathbb{P}}(u)$ so denoting the stochastic exponential by $\mathcal{E}(\cdot)$, we may write

$$\begin{aligned} v_u &= v_t \exp \left\{ \eta \tilde{W}_t^{\mathbb{P}}(u) + 2\nu C_H Z_t(u) \right\} \\ &= \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left(\eta \tilde{W}_t^{\mathbb{P}}(u) \right). \end{aligned} \quad (7)$$

- The conditional distribution of v_u depends on \mathcal{F}_t only through the variance forecasts $\mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t]$,
- To price options, one does not need to know \mathcal{F}_t , the entire history of the Brownian motion $W_s^{\mathbb{P}}$ for $s < t$.

Pricing under \mathbb{Q}

Our model under \mathbb{P} reads:

$$v_u = \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left(\eta \tilde{W}_t^{\mathbb{P}}(u) \right). \quad (8)$$

Consider some general change of measure

$$dW_s^{\mathbb{P}} = dW_s^{\mathbb{Q}} + \lambda_s ds,$$

where $\{\lambda_s : s > t\}$ has a natural interpretation as the price of volatility risk. We may then rewrite (8) as

$$v_u = \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left(\eta \tilde{W}_t^{\mathbb{Q}}(u) \right) \exp \left\{ \eta \sqrt{2H} \int_t^u \frac{\lambda_s}{(u-s)^\gamma} ds \right\}.$$

- Although the conditional distribution of v_u under \mathbb{P} is lognormal, it will not be lognormal in general under \mathbb{Q} .
 - The upward sloping smile in VIX options means λ_s cannot be deterministic in this picture.

The rough Bergomi (rBergomi) model

Let's nevertheless consider the simplest change of measure

$$dW_s^{\mathbb{P}} = dW_s^{\mathbb{Q}} + \lambda(s) ds,$$

where $\lambda(s)$ is a deterministic function of s . Then from (34), we would have

$$\begin{aligned} v_u &= \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left(\eta \tilde{W}_t^{\mathbb{Q}}(u) \right) \exp \left\{ \eta \sqrt{2H} \int_t^u \frac{1}{(u-s)^{\gamma}} \lambda(s) ds \right\} \\ &= \xi_t(u) \mathcal{E} \left(\eta \tilde{W}_t^{\mathbb{Q}}(u) \right) \end{aligned} \quad (9)$$

where the forward variances $\xi_t(u) = \mathbb{E}^{\mathbb{Q}} [v_u | \mathcal{F}_t]$ are (at least in principle) tradable and observed in the market.

- $\xi_t(u)$ is the product of two terms:
 - $\mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t]$ which depends on the historical path $\{W_s, s < t\}$ of the Brownian motion
 - a term which depends on the price of risk $\lambda(s)$.

Features of the rough Bergomi model

- The rBergomi model is a non-Markovian generalization of the Bergomi model:

$$\mathbb{E}[v_u | \mathcal{F}_t] \neq \mathbb{E}[v_u | v_t].$$

- The rBergomi model is Markovian in the (infinite-dimensional) state vector $\mathbb{E}^{\mathbb{Q}}[v_u | \mathcal{F}_t] = \xi_t(u)$.
- We have achieved our aim of replacing the exponential kernels in the Bergomi model (3) with a power-law kernel.
 - We may therefore expect that the rBergomi model will generate a realistic term structure of ATM volatility skew.

The stock price process

- The observed anticorrelation between price moves and volatility moves may be modeled naturally by anticorrelating the Brownian motion W that drives the volatility process with the Brownian motion driving the price process.
- Thus

$$\frac{dS_t}{S_t} = \sqrt{v_t} dZ_t$$

with

$$dZ_t = \rho dW_t + \sqrt{1 - \rho^2} dW_t^\perp$$

where ρ is the correlation between volatility moves and price moves.

Hybrid simulation of BSS processes

- In [BFG16], we simulate the rBergomi model by generating paths of \tilde{W} and Z with the correct joint marginals using Cholesky decomposition.
 - This is very slow!
- The rBergomi variance process is a special case of a Brownian Semistationary (BSS) process.
- In [BLP17], Bennedsen et al. show how to simulate such processes more efficiently.
 - Their hybrid BSS scheme is much more efficient than the exact simulation described above.
 - An even more efficient scheme is presented in [MP17].
 - However, it is still not fast enough to enable efficient calibration of the Rough Bergomi model to the volatility surface.

Guessing rBergomi model parameters

- The rBergomi model has only three parameters: H , η and ρ .
- These parameters have very direct interpretations:
 - H controls the decay of ATM skew $\psi(\tau)$ for very short expirations
 - The product $\rho\eta$ sets the level of the ATM skew for longer expirations.
 - Keeping $\rho\eta$ constant but decreasing ρ (so as to make it more negative) pushes the minimum of each smile towards higher strikes.
- So we can guess parameters in practice.

SPX smiles in the rBergomi model

- In Figure 9, we show how well a rBergomi model simulation with guessed parameters fits the SPX option market as of August 14, 2013, one trading day before the third Friday expiration.
 - Options set at the open of August 16, 2013 so only one trading day left.
- rBergomi parameters were: $H = 0.05$, $\eta = 2.3$, $\rho = -0.9$.
 - Only three parameters to get a very good fit to the whole SPX volatility surface!
- Note in particular that the extreme short-dated smile is well reproduced by the rBergomi model.
 - There is no need to add jumps!

SPX smiles as of August 14, 2013

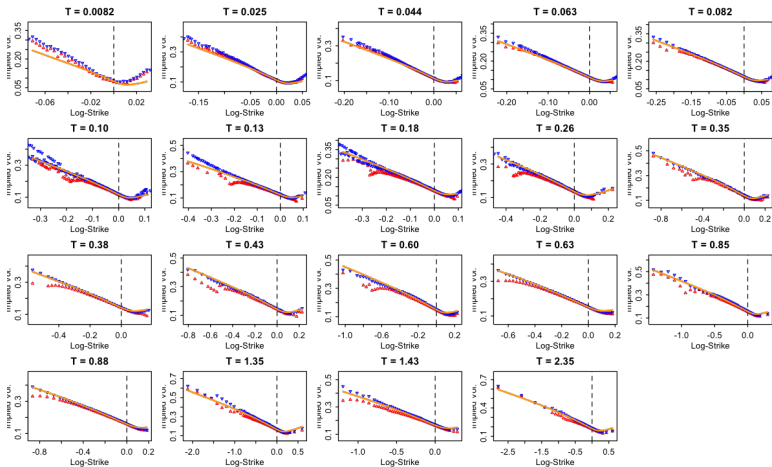


Figure 8: Red and blue points represent bid and offer SPX implied volatilities; orange smiles are from the rBergomi simulation.

The one-month SPX smile as of August 14, 2013




Figure 9: Red and blue points represent bid and offer SPX implied volatilities; the orange smiles is from the rBergomi simulation.

ATM volatilities and skews

In Figures 10 and 11, we see just how well the rBergomi model can match empirical ATM vols and skews. Recall also that the parameters we used are just guesses!

Term structure of ATM vol as of August 14, 2013




Figure 10: Blue points are empirical ATM volatilities; green points are from the rBergomi simulation. The two match very closely, as they should.

Term structure of ATM skew as of August 14, 2013




Figure 11: Blue points are empirical skews; the red line is from the rBergomi simulation.

Calibration

- As mentioned earlier, computation in the rBergomi model is challenging so calibration is not easy.
 - Though as noted earlier, we can easily guess parameters.
- We have investigated a number of approaches to calibration
 - Asymptotic expansions
 - Chebyshev interpolation
 - Moment matching
- So far, we cannot claim to have had real success with any of these approaches.
- A recent development shows some promise however...

The Alòs Itô decomposition formula

Following Elisa Alòs in [Alò12], let $X_t = \log S_t/K$ and consider the price process

$$dX_t = \sigma_t dZ_t - \frac{1}{2} \sigma_t^2 dt.$$


Now let $H(x, w)$ be some function that solves the Black-Scholes equation.

- Specifically,

$$-\partial_w H(x, w) + \frac{1}{2} (\partial_{xx} - \partial_x) H(x, w) = 0$$

which is of course the gamma-vega relationship.

- Note in particular that ∂_x and ∂_w commute when applied to a solution of the Black-Scholes equation.



Freezing derivatives

Freezing the derivatives in the Alòs Itô decomposition formula (10) gives us the approximation

$$\begin{aligned}\mathbb{E}[H_T | \mathcal{F}_t] &\approx H_t + \mathbb{E} \left[\int_t^T d\langle X, M \rangle_s \middle| \mathcal{F}_t \right] \partial_{xw} H_t \\ &\quad + \frac{1}{2} \mathbb{E} \left[\int_t^T d\langle M, M \rangle_s \middle| \mathcal{F}_t \right] \partial_{ww} H_t \\ &= H_t + (X \diamond M)_t(T) \cdot H_t + \frac{1}{2} (M \diamond M)_t(T) \cdot H_t.\end{aligned}$$

Remark

The essence of the Exponentiation Theorem is that we may express $\mathbb{E}[H_T | \mathcal{F}_t]$ as an exact expansion consisting of infinitely many terms, with derivatives in each such term frozen.

Diamond and dot notation

Let A_t and B_t be stochastic processes (some combinations of X_t and M_t). Then

$$(A \diamond B)_t(T) = \mathbb{E} \left[\int_t^T d\langle A, B \rangle_s \middle| \mathcal{F}_t \right].$$

When $(A \diamond B)_t(T)$ appears before some solution H_t of the Black-Scholes equation, the dot \cdot is to be understood as representing the action of ∂_x and ∂_w applied to H_t .

So for example

$$(X \diamond M)_t(T) \cdot H_t = \mathbb{E} \left[\int_t^T d\langle X, M \rangle_s \middle| \mathcal{F}_t \right] \partial_{xw} H_t$$

and so on.

Trees

- Terms such as $(X \diamond M)$, $(M \diamond M)$ and $X \diamond (X \diamond M)$ are naturally indexed by trees, each of whose leaves corresponds to either X or M .
- We end up with diamond trees reminiscent of Feynman diagrams, with analogous rules.

The first few forests

Applying this definition to compute the first few terms, we obtain


$$\mathbb{F}_0 = M$$

$$\mathbb{F}_1 = X \diamond \mathbb{F}_0 = (X \diamond M)$$

$$\mathbb{F}_2 = \frac{1}{2}(\mathbb{F}_0 \diamond \mathbb{F}_0) + X \diamond \mathbb{F}_1 = \frac{1}{2}(M \diamond M) + X \diamond (X \diamond M)$$



$$\begin{aligned} \mathbb{F}_3 &= (\mathbb{F}_0 \diamond \mathbb{F}_1) + X \diamond \mathbb{F}_2 \\ &= M \diamond (X \diamond M) + \frac{1}{2} X \diamond (M \diamond M) + X \diamond (X \diamond (X \diamond M)) \end{aligned}$$

The first forest $\mathbb{F}_1 = X \diamond M$



The second forest \mathbb{F}_2

$$\mathbb{F}_2 = \frac{1}{2}(M \diamond M) + X \diamond (X \diamond M)$$



Simple diamond rules

- For $k > 0$, the k th forest \mathbb{F}_k contains all trees with $k + 2$ leaves where X is counted as a single leaf, and M as a double leaf.
- Prefactor computation:
 - Work from the bottom up.
 - If child subtrees immediately below a diamond node are identical, carry a multiplicative factor of $\frac{1}{2}$.

Then from the Exponentiation Theorem, we have the following lemma.

Lemma

Let

$$\varphi_t^T(a) = \mathbb{E} \left[e^{i a X_T} \middle| \mathcal{F}_t \right]$$

be the characteristic function of the log stock price. Then

$$\varphi_t^T(a) = e^{\sum_{k=1}^{\infty} \tilde{\mathbb{F}}_k(a)} \Phi_t^T(a).$$

Corollary

The cumulant generating function (CGF) is given by

$$\psi_t^T(a) = \log \varphi_t^T(a) = i a X_t - \frac{1}{2} a (a + i) w_t(T) + \sum_{k=1}^{\infty} \tilde{\mathbb{F}}_k(a).$$

Leverage swap from the smile

Let

$$d_{\pm}(k) = \frac{-k}{\sigma_{\text{BS}}(k, T)\sqrt{T}} \pm \frac{\sigma_{\text{BS}}(k, T)\sqrt{T}}{2}$$

and following Fukasawa in [Fuk12], denote the inverse functions by $g_{\pm}(z) = d_{\pm}^{-1}(z)$. Further define

$$\sigma_{\pm}(z) = \sigma_{\text{BS}}(g_{\pm}(z), T)\sqrt{T}.$$




- Then the variance swap may be estimated from the smile using

$$w_t(T) = \int_{\mathbb{R}} dz N'(z) \sigma_-^2(z).$$

- And the gamma swap may be estimated using

$$\mathcal{G}_t(T) = \int_{\mathbb{R}} dz N'(z) \sigma_+^2(z).$$

- Recall that $\mathcal{L}_t(T) = \mathcal{G}_t(T) - w_t(T)$.
- We may thus compute leverage swaps in any stochastic volatility model (in principle) and also estimate leverage swaps from the implied volatility surface.
 - Model calibration is possible!



The rough Heston model

In the zero mean reversion limit, the rough Heston model of [ER16] may be written as

$$\frac{dS_t}{S_t} = \sqrt{v_t} \left\{ \rho dW_t + \sqrt{1 - \rho^2} dW_t^\perp \right\} = \sqrt{v_t} dZ_t$$

with

$$v_u = \xi_t(u) + \frac{\nu}{\Gamma(\alpha)} \int_t^u \frac{\sqrt{v_s}}{(u-s)^\gamma} dW_s, \quad u \geq t$$

where $\xi_t(u) = \mathbb{E}[v_u | \mathcal{F}_t]$ is the forward variance curve, $\gamma = \frac{1}{2} - H$ and $\alpha = 1 - \gamma = H + \frac{1}{2}$.

The rough Heston model in forward variance form

In forward variance form,

$$d\xi_t(u) = \frac{\nu}{\Gamma(\alpha)} \frac{\sqrt{v_t}}{(u-t)^\alpha} dW_t. \quad (13)$$

Remark

(13) is a natural fractional generalization of the classical Heston model which reads, in forward variance form [BG12],

$$d\xi_t(u) = \nu \sqrt{v_t} e^{-\kappa(u-t)} dW_t.$$

Computation of autocovariance functionals

Apart from \mathcal{F}_t measurable terms (abbreviated as 'drift'), we have

$$\begin{aligned}dX_t &= \sqrt{v_t} dZ_t + \text{drift} \\dM_t &= \int_t^T d\xi_t(u) du \\&= \frac{\nu}{\Gamma(\alpha)} \sqrt{v_t} \left(\int_t^T \frac{du}{(u-t)^\gamma} \right) dW_t \\&= \frac{\nu(T-t)^\alpha}{\Gamma(1+\alpha)} \sqrt{v_t} dW_t.\end{aligned}$$

Higher order forests

Define for $j \geq 0$

$$I_t^{(j)}(T) := \int_t^T ds \xi_t(s) (T - s)^{j\alpha}.$$

Then

$$\begin{aligned} dI_s^{(j)}(T) &= \int_s^T du d\xi_s(u) (T - u)^{j\alpha} + \text{drift terms} \\ &= \frac{\nu \sqrt{v_s}}{\Gamma(\alpha)} dW_s \int_s^T \frac{(T - u)^{j\alpha}}{(u - s)^\gamma} du + \text{drift terms} \\ &= \frac{\Gamma(1 + j\alpha)}{\Gamma(1 + (j + 1)\alpha)} \nu \sqrt{v_s} (T - s)^{(j+1)\alpha} dW_s + \text{drift terms.} \end{aligned}$$

With this notation,

$$(X \diamond M)_t(T) = \frac{\rho \nu}{\Gamma(1 + \alpha)} I_t^{(1)}(T).$$

The second order forest

There are two trees in \mathbb{F}_2 :

$$\begin{aligned}
 (M \diamond M)_t(T) &= \mathbb{E} \left[\int_t^T d\langle M, M \rangle_s \middle| \mathcal{F}_t \right] \\
 &= \frac{\nu^2}{\Gamma(1 + \alpha)^2} \int_t^T \xi_t(s) (T - s)^{2\alpha} ds \\
 &= \frac{\nu^2}{\Gamma(1 + \alpha)^2} I_t^{(2)}(T)
 \end{aligned}$$

and

$$\begin{aligned}
 (X \diamond (X \diamond M))_t(T) &= \frac{\rho \nu}{\Gamma(1 + \alpha)} \mathbb{E} \left[\int_t^T d\langle X, I^{(1)} \rangle_s \middle| \mathcal{F}_t \right] \\
 &= \frac{\rho^2 \nu^2}{\Gamma(1 + 2\alpha)} I_t^{(2)}(T).
 \end{aligned}$$

The third order forest

Continuing to the forest \mathbb{F}_3 , we have the following.


$$(M \diamond (X \diamond M))_t(T) = \frac{\rho \nu^3}{\Gamma(1 + \alpha) \Gamma(1 + 2\alpha)} I_t^{(3)}(T)$$

$$(X \diamond (X \diamond (X \diamond M)))_t(T) = \frac{\rho^3 \nu^3}{\Gamma(1 + 3\alpha)} I_t^{(3)}(T)$$

$$(X \diamond (M \diamond M))_t(T) = \frac{\rho \nu^3 \Gamma(1 + 2\alpha)}{\Gamma(1 + \alpha)^2 \Gamma(1 + 3\alpha)} I_t^{(3)}(T).$$

In particular, we easily identify the pattern

$$(X \diamond)^k M_t(T) = \frac{(\rho \nu)^k}{\Gamma(1 + k\alpha)} I_t^{(k)}(T).$$



The leverage swap under rough Heston

Using (12), we have

$$\begin{aligned} \mathcal{L}_t(T) &= \sum_{k=1}^{\infty} (X_{\diamond})^k M_t(T) \\ &= \sum_{k=1}^{\infty} \frac{(\rho \nu)^k}{\Gamma(1+k\alpha)} \int_t^T du \xi_t(u) (T-u)^{k\alpha} \\ &= \int_t^T du \xi_t(u) \{E_{\alpha}(\rho \nu (T-u)^{\alpha}) - 1\} \end{aligned} \quad (14)$$

where $E_{\alpha}(\cdot)$ denotes the Mittag-Leffler function.

An explicit expression for the leverage swap!

The normalized leverage swap

Given the form of equation (14), it is natural to normalize the leverage swap by the variance swap. We therefore define

$$L_t(T) = \frac{\mathcal{L}_t(T)}{w_t(T)}. \quad (15)$$

In the special case of the rough Heston model with a flat forward variance curve,

$$L_t(T) = E_{\alpha,2}(\rho \nu \tau^\alpha) - 1,$$

where $E_{\alpha,2}(\cdot)$ is a generalized Mittag-Leffler function. We further define an n th order approximation to $L_t(T)$ as

$$L_t^{(n)}(T) = \sum_{k=1}^n \frac{(\rho \nu \tau^\alpha)^k}{\Gamma(2 + k \alpha)}.$$

A numerical example

We now perform a numerical computation of the value of the leverage swap using the forest expansion in the rough Heston model with the following parameters, calibrated to the SPX options market as of April 24, 2017:

$$H = 0.0236; \quad \nu = 0.3266; \quad \rho = -0.6510.$$

The leverage swap under rough Heston

In Figure 12, we plot the normalized leverage swap $L_t(T)$ and successive approximations $L_t^{(n)}(T)$ to it as a function of τ .




Figure 12: Successive approximations to the (absolute value of) the normalized rough Heston leverage swap. The solid red line is the exact expression $L_t(T)$; $L_t^{(1)}(T)$, $L_t^{(2)}(T)$, and $L_t^{(3)}(T)$ are brown dashed, blue dotted and dark green dash-dotted lines respectively.

The leverage swap under rough Heston

We note that three terms are enough to get a very good approximation to the normalized leverage swap for all expirations traded in the listed market. Moreover, leverage swaps are straightforward to estimate from volatility smiles.

Remark

In practice, (15) can be used for very fast and efficient calibration of the three parameters of the rough Heston model by minimizing the distance between model and empirical normalized leverage swap estimates.

Leverage estimates using VolaDynamics




Figure 13: Leverage estimates using the VolaDynamics curves C13PM (red) and C14PM (blue) and their respective rough Heston fits as of 24-Apr-2017. See <https://voladynamics.com>.

- With a good volatility surface parameterization, the term structure of normalized leverage can be robustly estimated.

Summary

- We uncovered a remarkable monofractal scaling relationship in historical volatility which now appears to be universal.
- This leads to a natural non-Markovian stochastic volatility model under \mathbb{P} .
- The resulting volatility forecast beats existing alternatives.
- The simplest specification of $\frac{dQ}{dP}$ gives a non-Markovian generalization of the Bergomi model.
 - The history of the Brownian motion $\{W_s, s < t\}$ required for pricing is encoded in the forward variance curve, which is observed in the market.
- This model fits the observed volatility surface surprisingly well with very few parameters.
- Efficient calibration of the model to the volatility surface is now within reach.

References



Elisa Alòs, Jorge A León, and Josep Vives.

On the short-time behavior of the implied volatility for jump-diffusion models with stochastic volatility.
Finance and Stochastics, 11(4):571–589, 2007.



Elisa Alòs.

A decomposition formula for option prices in the Heston model and applications to option pricing approximation.
Finance and Stochastics, 16(3):403–422, 2012.



Elisa Alòs, Jim Gatheral, and Radoš Radoičić.

Exponentiation of conditional expectations under stochastic volatility.
SSRN, 2017.



Emmanuel Bacry and Jean François Muzy.

Log-infinitely divisible multifractal processes.
Communications in Mathematical Physics, 236(3):449–475, 2003.



Emmanuel Bacry and Jean-François Muzy.

Hawkes model for price and trades high-frequency dynamics.
Quantitative Finance, 14(7):1147–1166, 2014.



Christian Bayer, Peter Friz, and Jim Gatheral.

Pricing under rough volatility.
Quantitative Finance, 16(6):887–904, 2016.



Lorenzo Bergomi and Julien Guyon.

Stochastic volatility's orderly smiles.

Risk May, pages 60–66, 2012.



Mikkel Bennedsen, Asger Lunde, and Mikko S Pakkanen.

Hybrid scheme for Brownian semistationary processes.

Finance and Stochastics, 21(4): 931-965, 2017.



Mikkel Bennedsen, Asger Lunde, and Mikko S Pakkanen.

Decoupling the short-and long-term behavior of stochastic volatility.

Available at SSRN 2846756, 2016.



Omar El Euch and Mathieu Rosenbaum.

The characteristic function of rough Heston models.

arXiv preprint arXiv:1609.02108, 2016.



Masaaki Fukasawa.

Asymptotic analysis for stochastic volatility: Martingale expansion.

Finance and Stochastics, 15(4):635–654, 2011.



Masaaki Fukasawa.

The normalizing transformation of the implied volatility smile.

Mathematical Finance, 22(4):753–762, 2012.

