10 years of rough volatility: A current perspective

Jim Gatheral

Baruch COLLEGE
WEISSMAN SCHOOL OF ARTS AND SCIENCES
Master of Financial Engineering

XXV Workshop on Quantitative Finance, Bologna
April 11, 2024

Outline of this talk

- The shape of the volatility surface
  - Scaling of implied volatility smiles
- Monofractal scaling of historical volatility
- The Rough Fractional Stochastic Volatility (RFSV) model
- Estimating the roughness parameter H
- Forecasting realized variance
- Rough volatility models (under $\mathbb{Q}$)
- Computation
  - The hybrid scheme
  - Markovian approximations
  - Rational approximation of rough Heston
- What's next?

The SPX volatility surface as of 15-Feb-2023

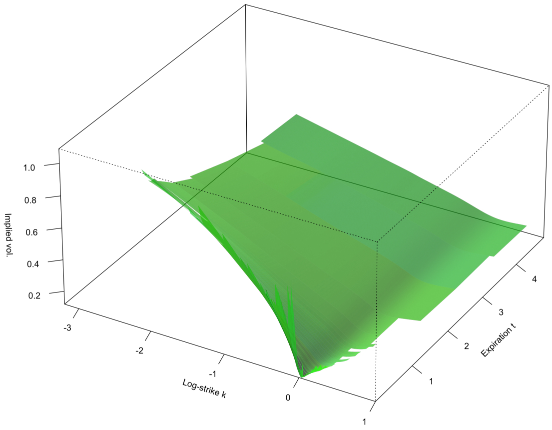

Figure 1: The SPX volatility surface as of 15-Feb-2023.[1]

[1] Data from OptionMetrics via WRDS.

Remarks on Figure 1

- Figure 1 is a slightly smoothed plot of estimated mid volatilities, not a fit!
  - There were 48 expirations and 6,749 put/call option pairs with non-zero bids as of the close on 15-Feb-2023.
- Notice how smooth this volatility surface is!
  - Bumps or dips would be tradable.

- Although the level and orientation of the volatility surface changes over time, it is a stylized fact that its rough shape stays very much the same.
  - The surface as of 15-Feb-2023 is typical.

SPX volatility smiles as of 15-Feb-2023

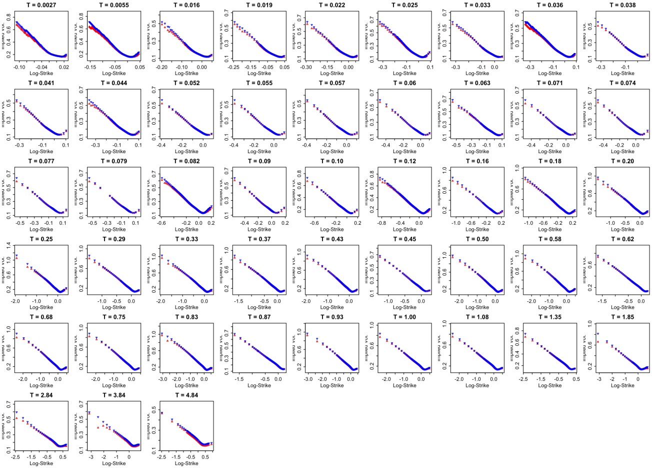

Figure 2: SPX volatility smiles as of 15-Feb-2023.

Interpreting a single smile

- We could say that the volatility smile (at least in equity markets) reflects two basic observations:
  - Volatility tends to increase when the underlying price falls,
    - hence the negative skew.
  - We don't know in advance what realized volatility will be,
    - hence implied volatility is increasing in the wings.
- It's implicit in the above that more or less any model that is consistent with these two observations will be able to fit one given smile.
  - Fitting two or more smiles simultaneously is much harder.

Term structure of at-the-money skew

- Given one smile for a fixed expiration, little can be said about the process generating it.
- In contrast, the dependence of the smile on time to expiration is intimately related to the underlying dynamics.
  - In particular model estimates of the term structure of ATM volatility skew defined as

$$\psi(\tau) := \left. \frac{\partial}{\partial k} \sigma_{\mathrm{BS}}(k, \tau) \right|_{k=0}$$

are very sensitive to the choice of volatility dynamics in a stochastic volatility model.

Term structure of SPX ATM skew as of 15-Feb-2023

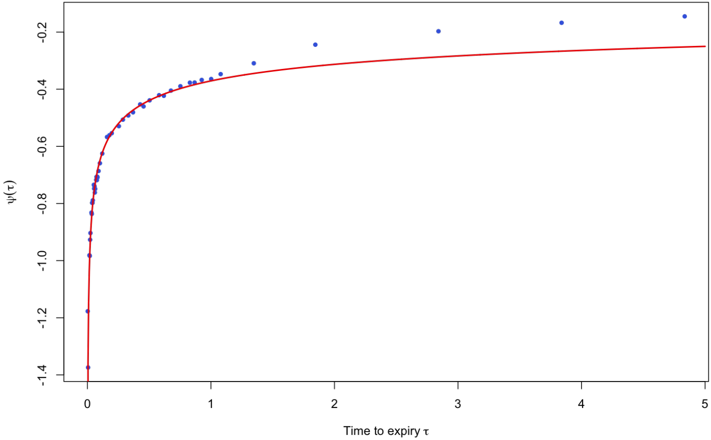

Figure 3: Term structure of ATM skew as of 15-Feb-2023, with power law fit $\tau^{-0.24}$ superimposed in red.

Stochastic volatility models

- A generic stochastic volatility model takes the form

$$\frac{dS_t}{S_t} = \sqrt{V_t}\, dZ_t$$
$$V_t = \int_{-\infty}^{t} F(\Omega_s)\, dW_s,$$

where $V_t\, dt = d\,\langle \log S \rangle_t$, F is some function, and Ω_t is the natural filtration generated by Z and W.

Fractional stochastic volatility models

- Non-Markovian models of the form

$$V_t = V_0 \exp\left\{\eta \int_0^t \frac{dW_s}{(t-s)^\gamma} + \text{ drift }\right\}$$

were shown by Alòs et al. in [ALV07] and then by Fukasawa in [Fuk11] to generate a short-dated ATM skew of the form

$$\psi(\tau) \sim \tau^{-\gamma}$$

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

  - Such models, where the kernel decays as a power-law for small times, are called *rough volatility* models.

- The typical power-law behavior of the skew term structure for short times is one of the motivations for rough volatility models.

Skew term structure is not always power-law

- [GA23] show that a power-law fits the *average* skew term structure poorly for very short dates.
- [DDS23] show that skew term structure is typically a combination of two power-laws.

- We confirm in Figure 4, that the term structure of skew is not always power-law.
  - On 27-Dec-2022, the skew term structure is not even monotonic!
- We further confirm in Figure 5 that the skew term structure looks like a combination of two power-laws, at least on 15-Feb-2023, consistent with [DDS23].

Skew term structure is not always power-law

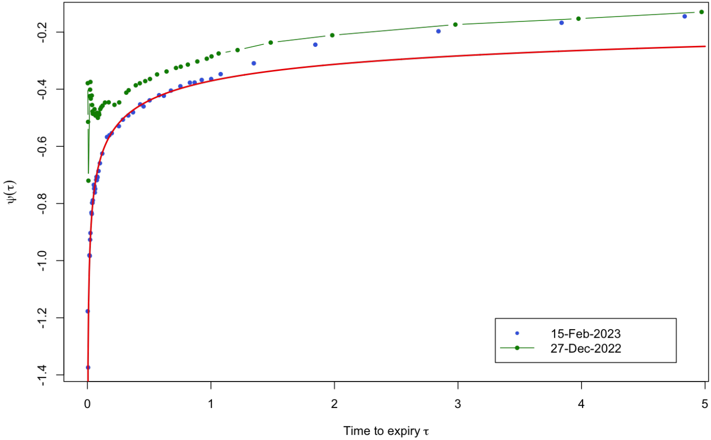

Figure 4: ATM skew term structure on two different dates. On 27-Dec-2022, the skew term structure is not even monotonic!

Log plot of skew term structure

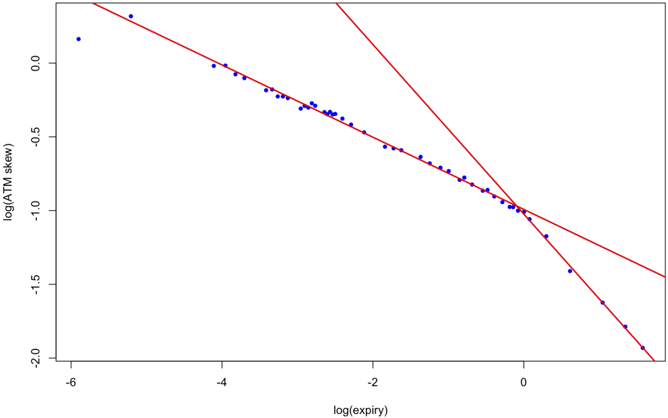

Figure 5: The skew term structure on 15-Feb-2023 looks like a superposition of two power-laws.

Scaling of total variance

- The rough SABR formula of [FG22] suggests that we should have

$$\frac{w(k,\tau)}{w(0,\tau)} \approx f\left(\tau^{-\gamma}\,\frac{k}{\Sigma_{BS}(0)}\right).$$

  - Roughly speaking, total variance curves should scale as a power-law.
- Figure 6 does suggest close-to-power-law scaling, even in the 27-Dec-2022 case.

Scaling of total variance

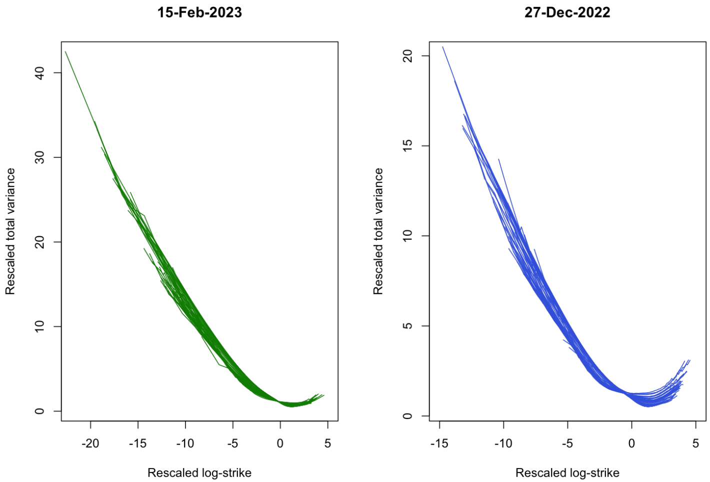

Figure 6: ATM skew term structure on two different dates. On 27-Dec-2022, the skew term structure is not even monotonic!

Fractional stochastic volatility models

- This simple scaling of volatility smiles suggests that rough volatility models should be consistent with option prices.
  - Despite that the term structure of skew is not always power-law.
- Were the instantaneous variance to follow something like

$$V_t = V_0 \exp\left\{\eta \int_0^t \frac{dW_s}{(t-s)^\gamma} + \text{ drift } \right\},$$

  the time series of $\log V_t$ should also have simple scaling properties.

- Let's check ...

Power-law scaling of the historical volatility process

- The Oxford-Man Institute of Quantitative Finance used to make historical realized variance (RV) estimates freely available.
  - Unfortunately, no longer. The last date in my dataset is 06/28/2022.
- Using daily RV estimates as proxies for instantaneous variance, we may investigate the time series properties of V_t empirically.

History of SPX realized variance

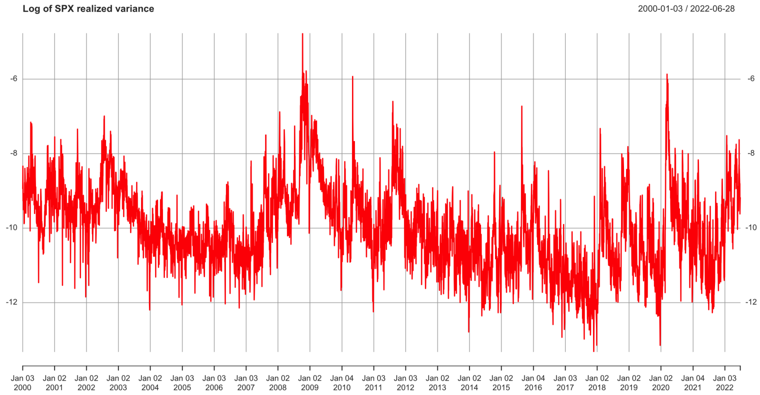

Figure 7: (Log) realized kernel estimates of SPX realized variance.

The smoothness of the volatility process

- For $q \geq 0$, we define the qth sample moment of differences of log-volatility at a given lag Δ[2]:

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

- For example

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

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

[2] $\langle \cdot \rangle$ denotes the sample average.

Scaling of $m(q, \Delta)$ with lag Δ

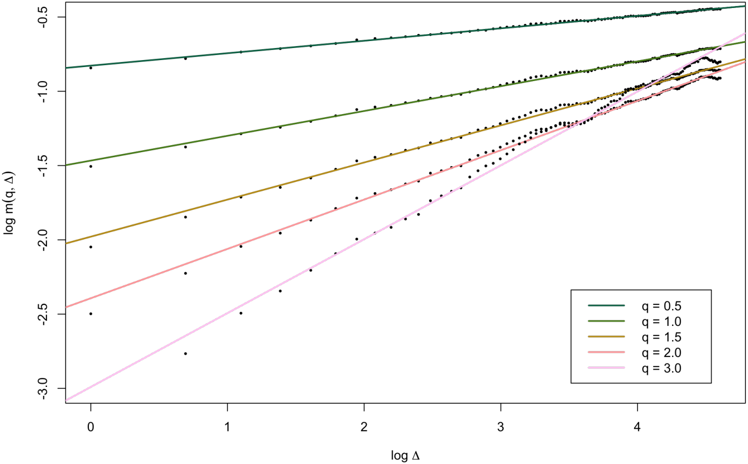

Figure 8: $\log m(q, \Delta)$ as a function of $\log \Delta$, SPX.

Scaling of ζ_q with q

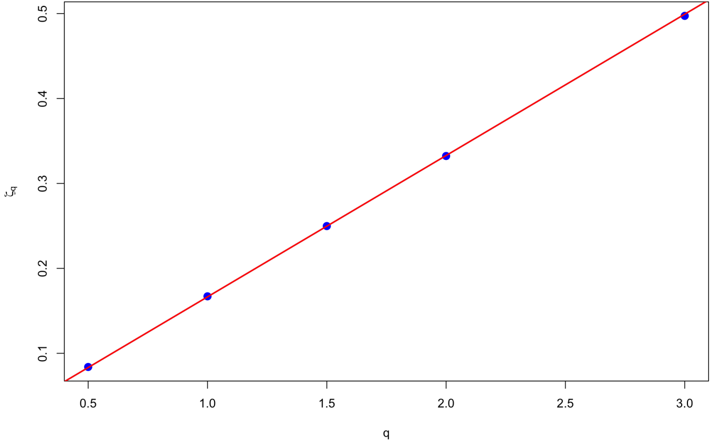

Figure 9: Scaling of ζ_q with q.

Monofractal scaling result

- From the log-log plot Figure 8, we see that for each q, $m(q, \Delta) \propto \Delta^{\zeta_q}$.
- And from Figure 9 the monofractal scaling relationship

$$\zeta_q = q\,H$$

with $H \approx 0.17$.

  - 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.
- It is easily checked that H is not a constant, but varies with time.

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 [BLP22] 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:

$$\log \sigma_{t+\Delta} - \log \sigma_t = \nu \left(W^H_{t+\Delta} - W^H_t \right) \tag{1}$$

where W^H is fractional Brownian motion.
  - If σ were continuous, distributions of $log \sigma$ were really Gaussian, and if H were constant, this model would be unique!
- In [GJR18], we refer to a stationary version of (1) 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

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

  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.

More sophisticated estimators of H

- Numerous authors have pointed out that the estimates of H by linear regression in [GJR18] make sense only if estimation error is not too high.
  - A semimartingale volatility process with substantial estimation error would yield spuriously low estimates of H.
  - Some authors have *even* suggested that volatility may not be rough!
    - Easily rejected by examining the magnitude of ν.
- More sophisticated estimators of H include
  - The ACF estimator of [BLP22]
  - The Whittle estimator of [FT19]
  - The GMM estimator of [BCPV23]
  - The TDML estimator of [WXY223]
- All of these authors conclude that volatility of SPX is indeed rough.

Graph of ACF-estimated H

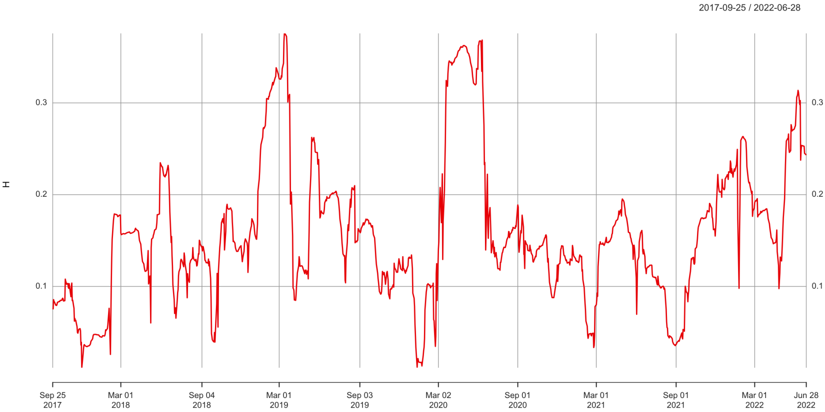

Figure 10: Plot of ACF estimates of H using 61 day windows. H is not constant!

The forecast formula

- In the RFSV model (1), $\log V_t \approx 2\,\nu\,W_t^H + C$ for some constant C.
- [NP00] show that $W^H_{t+\Delta}$ is conditionally Gaussian with conditional expectation

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

and conditional variance

$$\mathrm{Var}[W^H_{t+\Delta}|\mathcal{F}_t] = c\,\Delta^{2H}.$$

where

$$c = \frac{\Gamma(3/2-H)}{\Gamma(H+1/2)\,\Gamma(2-2H)}.$$

The forecast formula

- Thus, we obtain

Variance forecast formula

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

where

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

- [BLP22] 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

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

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

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

- Note that $\mathbb{E}^{\mathbb{P}}\left[M_t(u) \middle| \mathcal{F}_t \right] = 0$ 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 $s < t$!

The variance process under $\mathbb{P}$

Exponentiating (2), we get

$$
\begin{aligned}
V_u &= V_t \exp\left\{2\,\nu\,C_H\,M_t(u) + 2\,\nu\,C_H\,Z_t(u)\right\} \\
&= \mathbb{E}^{\mathbb{P}}\left[\,V_u \middle|\, \mathcal{F}_t\right] \mathcal{E}\left(2\,\nu\,C_H \int_t^u \frac{dW_s^{\mathbb{P}}}{(u-s)^{\gamma}}\right).
\end{aligned}
$$

- The conditional distribution of V_u depends on $\mathcal{F}_t$ only through the variance forecasts $\mathbb{E}^{\mathbb{P}}\left[\,V_u \middle|\, \mathcal{F}_t\right]$.
  - These variance forecasts depend explicitly on the history of the variance process.
  - Rough volatility models are (in principle) path-dependent!
- Rough volatility gives us a natural connection between $\mathbb{P}$ and $\mathbb{Q}$.
  - We can forecast the volatility surface using historical data!

Pricing under $\mathbb{Q}$

Our model under $\mathbb{P}$ reads:

$$V_u = \mathbb{E}^{\mathbb{P}}\left[V_u \middle| \mathcal{F}_t \right] \mathcal{E}\left(\tilde{\eta} \int_t^u \frac{dW_s^{\mathbb{P}}}{(u-s)^{\gamma}} \right), \tag{3}$$

where $\tilde{\eta} := 2\,\nu\, C_H$. 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 (3) as

$$V_u = \mathbb{E}^{\mathbb{P}}\left[V_u \middle| \mathcal{F}_t \right] \mathcal{E}\left(\tilde{\eta} \int_t^u \frac{dW_s^{\mathbb{Q}}}{(u-s)^{\gamma}} \right) \exp\left\{ \tilde{\eta} \int_t^u \frac{\lambda_s}{(u-s)^{\gamma}}\, ds \right\}. \tag{4}$$

- 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 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 (4), we would have

$$\begin{aligned} V_u &= \mathbb{E}^{\mathbb{P}}\left[V_u \middle| \mathcal{F}_t \right] \mathcal{E}\left(\tilde{\eta} \int_t^u \frac{dW_s^{\mathbb{Q}}}{(u-s)^{\gamma}} \right) \exp\left\{ \tilde{\eta} \int_t^u \frac{\lambda_s}{(u-s)^{\gamma}}\, ds \right\} \\ &= \xi_t(u)\, \mathcal{E}\left(\tilde{\eta} \int_t^u \frac{dW_s^{\mathbb{Q}}}{(u-s)^{\gamma}} \right). \end{aligned}$$

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

- $\xi_t(u)$ is the product of two terms:
  - $\mathbb{E}^{\mathbb{P}}\left[V_u \middle| \mathcal{F}_t \right]$ 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 rough Bergomi model is a non-Markovian generalization of the Bergomi model:

$$\mathbb{E}\left[V_u | \mathcal{F}_t \right] \neq \mathbb{E}[V_u | V_t].$$

- The rough Bergomi model is Markovian in the (infinite-dimensional) state vector $\mathbb{E}^{\mathbb{Q}}\left[V_u | \mathcal{F}_t \right] = \xi_t(u)$.
- From [ALV07] and [Fuk11], we expect the rough Bergomi model to generate a realistic term structure of ATM volatility skew.
  - It does!

The rough Heston model

- Rosenbaum et al. [EFR16] consider a simple Hawkes process model of order flow 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 $\alpha \approx 0.6$).
- In a *tour de force*, El Euch and Rosenbaum [ER19] compute an expression for the characteristic function of the rough Heston model.
  - In terms of the solution of a Riccati fractional ODE.

In forward variance form

- Let $\alpha = H + \frac{1}{2}$.
- The rough Bergomi model reads

$$d\xi_t(u) = \kappa(u-t)\,\xi_t(u)\,dW_t,$$

with kernel $\kappa(\tau) = \tilde{\eta}\,\tau^{\alpha-1}$.

- The rough Heston model reads

$$d\xi_t(u) = \kappa(u-t)\,\sqrt{V_t}\,dW_t,$$

with kernel $\kappa(\tau) = \nu\,\tau^{\alpha-1}\,E_{\alpha,\alpha}(-\lambda\tau^{\alpha})$, where $E_{\alpha,\alpha}(\cdot)$ is the Mittag-Leffler function.

A generalization

- [GKR19] show how this can be generalized by defining the class of *affine forward variance* (AFV) models.
  - We show that a one-factor stochastic volatility model is affine if and only if the volatility process is square-root.
  - For each such model, one can compute the characteristic function by solving a Riccati?Volterra equation.
  - Hawkes processes are the discrete time analogs of square-root stochastic volatility processes.
- Moreover, we define the class of affine forward intensity (AFI) models.
  - Each such model looks like Mathieu's rough Heston model but with a different Hawkes kernel.
  - Each such model has a rough Heston-like stochastic volatility model as a limit if and only if the Hawkes kernel is nearly unstable.

The quadratic rough Heston model

- Though the rough Heston model has a nice microstructural justification, as for the classical Heston model, its dynamics are not reasonable.
  - At-the-money skew $\psi(\tau)$ decreases with at-the-money volatility.
  - VIX smiles are downward-sloping.
- Lognormal dynamics are more reasonable.
- Inspired by [BDB19, DJR19], [GJR20] proposed the quadratic rough Heston model.
  - Because it is roughly speaking "the square" of a Heston process, the QR Heston process is roughly lognormal.
  - The QR Heston model is explicitly path-dependent: Variance is a weighted function of the price path.
  - Moreover, it generates upward-sloping VIX smiles!

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 rough volatility models

- In [BFG16], we simulate the rough Bergomi model by generating paths of $\tilde{W}$ and Z with the correct joint marginals using Cholesky decomposition.
  - This is very slow!
- The rough Bergomi 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 version of the hybrid scheme (with variance reduction) is presented in [MP18].

Idea of the hybrid scheme

- In the affine algorithm of [Gat22],

$$V_n = \xi_n + \sum_{k=1}^{n} \int_{(k-1)\Delta}^{k\Delta} \kappa(n\Delta - s)\,\sqrt{V_s}\,dW_s =: \hat{\xi}_n + U_n,$$

where the $\mathcal{F}_{n-1}$-adapted variable $\hat{\xi}_n$ is given by

$$\begin{aligned}\hat{\xi}_n = \mathbb{E}\left[\,V_n \middle|\, \mathcal{F}_{n-1}\right] &= \xi_n + \sum_{k=1}^{n-1} \int_{(k-1)\Delta}^{k\Delta} \kappa(n\Delta - s)\,\sqrt{V_s}\,dW_s \\ &\approx \xi_n + \sum_{k=1}^{n-1} \kappa(n\Delta - s)\,\chi_k, \end{aligned} \tag{5}$$

and the

$$\chi_k = \int_{(k-1)\Delta}^{k\Delta} \sqrt{V_s}\,dW_s,$$

are the kth increments of the log-stock price process.

- The martingale increment U_n given by

$$U_n = \int_{(n-1)\Delta}^{n\Delta} \kappa(n\Delta - s)\sqrt{V_s}\,dW_s,$$

  is simulated such that its variance and also covariance with the χ_k, $k < n$ are well-approximated.
    - Hence the terminology *hybrid scheme*.
- The convolution term $\sum_{k=1}^{n-1} \kappa(n\Delta - s)\,\chi_k$ in (5) directly captures the non-Markovianity of the rough Heston model.
    - Increasing complexity and slowing down computation.
- In the rough Bergomi model, to which the hybrid scheme was first applied [BLP17], the equivalent convolution can be speeded up using FFT.

Markovian approximations

- [AE19] show how to approximate the power-law kernel by a sum of exponential kernels.
  - Many Markovian simulations are required to accurately approximate a rough volatility model, so simulation is slow.
  - On the other hand, as pointed out in [Abi19], lower order approximations may themselves be considered models that display rough-like behavior over a large range of timescales.
- The idea is

$$\tau^{\alpha-1} = \frac{1}{\Gamma(1-\alpha)} \int_{\mathbb{R}^+} e^{-\kappa\,\tau} \frac{d\kappa}{\kappa^{\alpha}} \approx \sum_{i=1}^{n} w_i\, e^{-\kappa_i\,\tau}.$$

- [BB23a, BB23b] substantially improve on [AE19] by employing geometric grid spacing and Gaussian quadrature.
  - According to the authors, their 7-point approximation does better than a 1024 point [AE19]-style approximation!

The rough Heston model

- The rough Heston model with Mittag-Leffler kernel is particularly tractable.
- Its characteristic function is expressed as the solution of a fractional ODE which can be solved numerically using:
  - The Adams scheme [DFF04]
  - The HQE (hybrid) scheme [Gat22]
  - Markovian approximations
- Alternatively, the rational approximations of [GR19, GR24] are very accurate.
  - These rational approximations are a large factor faster than numerical methods, enabling fast calibration to the volatility surface.
- The rough Heston solution is thus a perfect benchmark for numerical algorithms.

What's next?

- The skew-stickiness ratio (SSR)
  - We need to explain why empirically observed skew-stickiness is so much lower than that generated from stochastic volatility models.
- Better numerical techniques?
  - Maybe we can combine insights from the (path-dependent) hybrid scheme and state-of-the-art Markovian approximation?
- Better fitting but tractable rough volatility models?
  - That is probably expecting too much!
- Economic rational for rough volatility?
  - We already have some understanding of how rough dynamics can be generated from market microstructure models.
  - Maybe there is a more fundamental economic rationale?

Postscript

From [GJR18]:

> *It is of course plausible that other models are compatible with many of our observations. In fact, there are probably many ways to design a process so that most of our empirical results are reproduced (for example estimation errors when estimating volatility can be quite significant for some models, leading to downward biases in the measurement of the smoothness). However, what we show here is that we cannot find any evidence against the RFSV model. In statistical terms, the null hypothesis that the data generating process of the volatility is a RFSV model cannot be rejected based on our analysis.*

From the preface of [BFFGJR23]:

> *... any viable model describing the dynamics of volatility should exhibit a close resemblance to rough volatility.*

From [Fuk23]:

> *The pathwise roughness of volatility itself is, however, only secondary; volatility is only a hypothetical latent quantity in a diffusive scale, and its realization cannot be identified from finite data. Its roughness is a result of the distributional property of local self-similarity under a Brownian semimartingale framework. In this sense, the statement "volatility is rough" seems misleading and, indeed, has invoked some misunderstanding and confusion in the mathematical finance community.*

From [Fuk23] again:

> *A relevant scientific question is why volatility is rough (or more precisely, why the distributions of asset prices in the daily time scale have such universal properties that are well-explained by rough volatility). There are several studies that connect rough volatility and market microstructure dynamics. An economic reasoning is still absent. Research in this direction would require a deeper understanding of stochastic processes.*

Summary

- Approximate power-law scaling of volatility smiles suggests a scaling relationship for instantaneous variance.
- This leads us to uncover a remarkable monofractal scaling relationship in historical volatility which now appears to be universal.
  - The rough volatility paradigm.
- A natural non-Markovian (path-dependent) stochastic volatility model under $\mathbb{P}$ then follows.
- The resulting volatility forecast beats existing alternatives.
- Rough volatility models fit the observed volatility surface remarkably well with very few parameters.
- Rough volatility models offer consistent modeling of historical and implied volatility.

Further reading ...

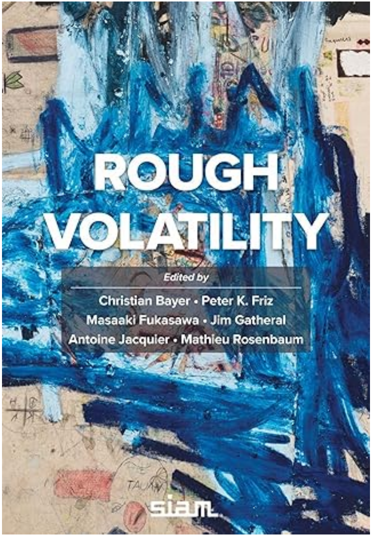

References

Eduardo Abi Jaber and Omar El Euch.
Multifactor approximation of rough volatility models.
SIAM Journal on Financial Mathematics, 10(2):309–349, 2019.

Eduardo Abi Jaber.
Lifting the Heston model.
Quantitative Finance, 19(12):1995–2013, 2019.

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.

Emmanuel Bacry and Jean François Muzy.
Log-infinitely divisible multifractal processes.
Communications in Mathematical Physics, 236(3):449–475, 2003.

Christian Bayer, Peter Friz, and Jim Gatheral.
Pricing under rough volatility.
Quantitative Finance, 16(6):887–904, 2016.

Christian Bayer and Simon Breneis.
Markovian approximations of stochastic Volterra equations with the fractional kernel.
Quantitative Finance, 23(1):53–70, 2023.

Christian Bayer and Simon Breneis.
Efficient option pricing in the rough Heston model using weak simulation schemes.
http://arxiv.org/abs/2310.04146, 2023.

Christian Bayer, Peter K Friz, Masaaki Fukasawa, Jim Gatheral, Antoine Jacquier, and Mathieu Rosenbaum.
Rough Volatility.
SIAM, 2023.

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.
Journal of Financial Econometrics, 20(5), 961–1006, 2022.

Pierre Blanc, Jonathan Donier and Jean-Philippe Bouchaud.
Quadratic Hawkes processes for financial prices.
Quantitative Finance, 17(2) 171–188, 2017.

Anine E. Bolko, Kim Christensen, Mikko S. Pakkanen and Bezirgen Veliyev.
A GMM approach to estimate the roughness of stochastic volatility.
Journal of Econometrics, 235(2) 745–778, 2023.

Carsten H. Chong and Victor Todorov.
The Fine Structure of Volatility Dynamics.
SSRN https://papers.ssrn.com/abstract=4660138, 2023.

Aditi Dandapani, Paul Jusselin, and Mathieu Rosenbaum
From quadratic Hawkes processes to super-Heston rough volatility models with Zumbach effect
Quantitative Finance, 21(8):1235–1247, 2021.

Jules Delemotte, Stefano De Marco and Florent Segonne,
Yet Another Analysis of the SP500 At-The-Money Skew: Crossover of Different Power-Law Behaviours
SSRN, https://papers.ssrn.com/abstract=4428407, 2023.

Kai Diethelm, Neville J. Ford, and Allan D. Freed
Detailed error analysis for a fractional Adams method.
Numerical Algorithms, 36(1) 31–52, 2004.

Omar El Euch, Masaaki Fukasawa, and Mathieu Rosenbaum.
The microstructural foundations of leverage effect and rough volatility.
Finance and Stochastics, 22(2) 241–280, 2018.

Omar El Euch and Mathieu Rosenbaum.
The characteristic function of rough Heston models.
Mathematical Finance, 29(1), 3–38, 2019.

Masaaki Fukasawa.
Asymptotic analysis for stochastic volatility: Martingale expansion.
Finance and Stochastics, 15(4):635–654, 2011.

Masaaki Fukasawa and Asuto Hirano
Refinement by reducing and reusing random numbers of the Hybrid scheme for Brownian semistationary processes.
Quantitative Finance, 21(7):1127–1146, 2021.

Masaaki Fukasawa and Tetsuya Takabatake.
Asymptotically efficient estimators for self-similar stationary Gaussian noises under high frequency observations.
Bernoulli, 25(3), 1870–1900, 2019.

Masaaki Fukasawa and Jim Gatheral.
A rough SABR formula.
Frontiers of Mathematical Finance, 1(1), 81–97, 2022.

Masaaki Fukasawa.
Wiener spiral for volatility modeling.
Teoriya veroyatnostey i yeye primeneniya, 68(3), 596–618, 2023.

Jim Gatheral.
The volatility surface: A practitioner's guide.
John Wiley & Sons, 2006.

Jim Gatheral, Paul Jusselin, and Mathieu Rosenbaum.
The quadratic rough Heston model and the joint S&P 500/VIX smile calibration problem.
Risk.net, May 2020.

Jim Gatheral, Thibault Jaisson, and Mathieu Rosenbaum.
Volatility is rough.
Quantitative Finance, 18(6):933–949, 2018.

Jim Gatheral and Martin Keller-Ressel
Affine forward variance models.
Finance and Stochastics, 23(3):501–533, 2019.

Jim Gatheral
Efficient simulation of affine forward variance models.
Risk.net, February 2022.

Jim Gatheral and Radoš Radoičić.
Rational approximation of the rough Heston solution.
International Journal of Theoretical and Applied Finance, 22(3):1950010–19, 2019.

Jim Gatheral and Radoš Radoičić.
A generalization of the rational rough Heston approximation.
Quantitative Finance, 24(2):329–335, 2024.

Julien Guyon and Mehdi El Amrani
Does the Term-Structure of the At-the-Money Skew Really Follow a Power Law?
Risk.net, August 2023.

Paul Jusselin and Mathieu Rosenbaum
No-arbitrage implies power-law market impact and rough volatility
Mathematical Finance, 30(4):1309–1336, 2020.

Ryan McCrickerd and Mikko S. Pakkanen
Turbocharging Monte Carlo pricing for the rough Bergomi model.
Quantitative Finance, 18(11) 1877–1886, 2018.

Carl J Nuzman and Vincent H Poor.
Linear estimation of self-similar processes via Lamperti's transformation.
Journal of Applied Probability, 37(2):429–452, 2000.

Xiaohu Wang, Weilin Xiao, and Jun Yu.
Modeling and forecasting realized volatility with the fractional Ornstein?Uhlenbeck process.
Journal of Econometrics, 232(2):389–415, 2023.