Volatility is Rough, Part 2: Pricing

Jim Gatheral
(joint work with Christian Bayer, Peter Friz,
Thibault Jaisson, Andrew Lesniewski, and Mathieu Rosenbaum)

Baruch College
The City University of New York

Workshop on Stochastic and Quantitative Finance,
Imperial College London,
Saturday November 29, 2014

Outline of this talk

- The volatility surface: Stylized facts
- The Rough Fractional Stochastic Volatility (RFSV) model
- The Rough Bergomi (rBergomi) model
- Change of measure

The SPX volatility surface as of 15-Sep-2005

- We begin by studying the SPX volatility surface as of the close on September 15, 2005.
  - Next morning is *triple witching* when options and futures set.
- We will plot the volatility smiles, superimposing an SVI fit.
  - SVI stands for "stochastic volatility inspired", a well-known parameterization of the volatility surface.
  - We show in [Gatheral and Jacquier] how to fit SVI to the volatility surface in such a way as to guarantee the absence of static arbitrage.
- We then interpolate the resulting SVI smiles to obtain and plot the whole volatility surface.

SPX volatility smiles as of 15-Sep-2005

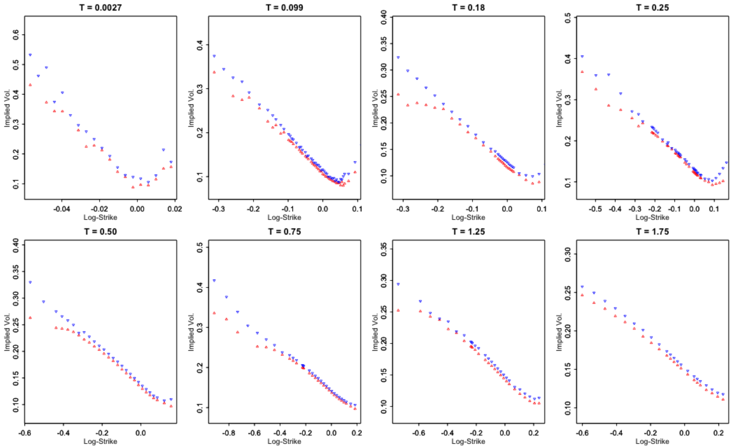

Figure 1: SPX volatility smiles as of 15-Sep-2005.

SPX volatility smiles as of 15-Sep-2005

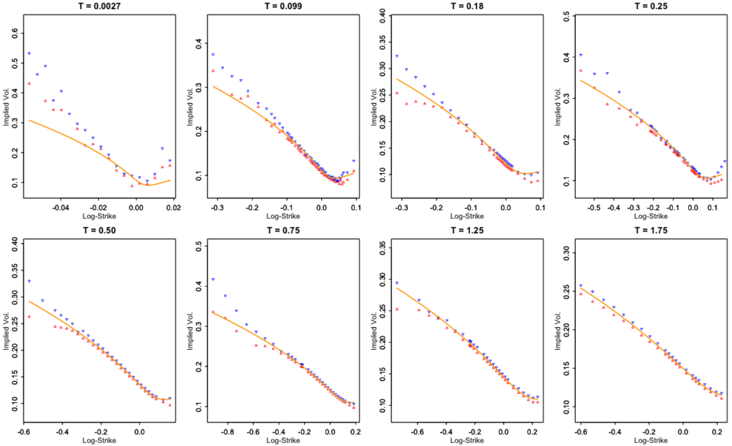

Figure 2: SVI fit superimposed on smiles.

The SPX volatility surface as of 15-Sep-2005

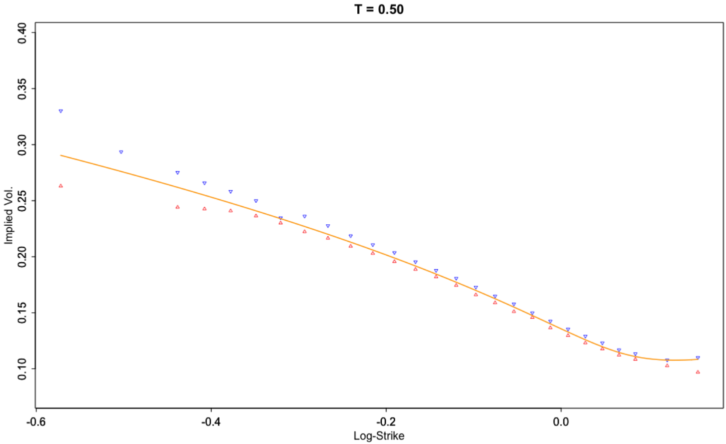

Figure 3: The March expiry smile as of 15-Sep-2005 – the SVI fit looks OK!

The SPX volatility surface as of 15-Sep-2005

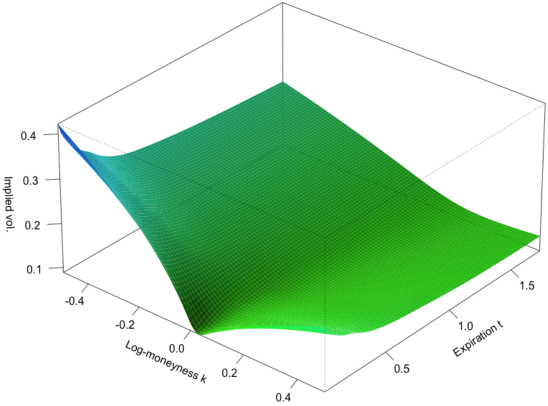

Figure 4: The SPX volatility surface as of 15-Sep-2005 (Figure 3.2 of The Volatility Surface).

Interpreting the 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.
    - Heston for example fits a maximum of two smiles simultaneously.
    - SABR can only fit one smile at a time.

Term structure of at-the-money skew

- What really distinguishes between models is how the generated smile depends on time to expiration.
    - In particular, their predictions for the term structure of ATM volatility skew defined as

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

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

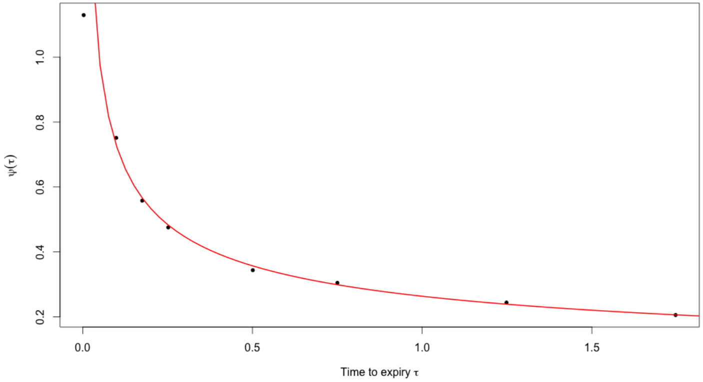

Figure 5: Term structure of ATM skew as of 15-Sep-2005, with power law fit $\tau^{-0.44}$ superimposed in red.

SPX volatility surfaces from 2005 to 2011

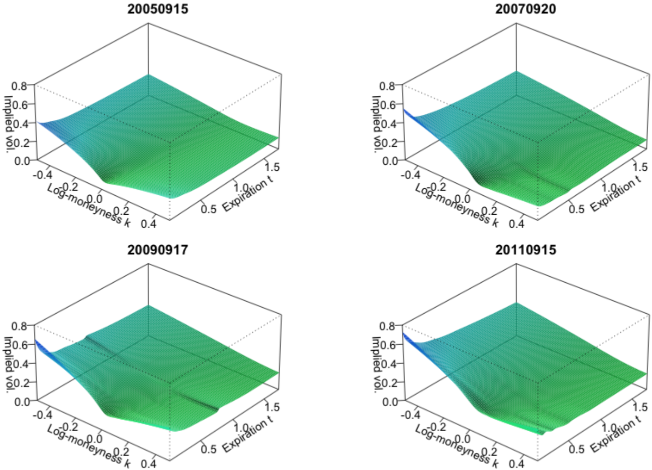

Figure 6: SPX volatility surfaces over the years as of the close before September SQ.

Observations

- We note that although the levels and orientations of the volatility surfaces change over time, their rough shape stays very much the same.

- Let's now look at the term structure of ATM skew on these dates.

Term structure of SPX ATM skew as over the years

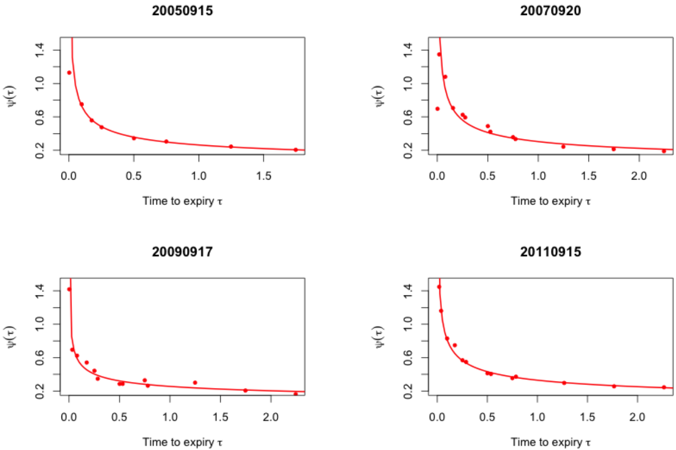

Figure 7: SPX ATM skew over the years as of the close before September SQ. Power-laws fits are superimposed.

Conclusion

- The shape of the volatility surface seems to be more-or-less stable.
  - 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 $\alpha \in (0.3, 0.5)$.

Conventional 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.

Bergomi Guyon

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

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

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

$$C^{x\,\xi} = \int_0^T dt \int_t^T du\, \frac{\mathbb{E}\left[dx_t\, d\xi_t(u)\right]}{dt}. \qquad (2)$$

- Thus, given a stochastic model, defined in terms of an SDE, we can easily (at least in principle) compute this smile approximation.

Connecting the time series with options prices

- Suppose for a moment that the pricing measure $\mathbb{Q}$ is the same as the historical (or physical) measure $\mathbb{P}$.
- Then equation (2) also connects the prices of options with statistics of the historical time series of volatility.

The Bergomi model

- The n-factor Bergomi variance curve model reads:

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

- To achieve a decent fit to the observed volatility surface, and to control the forward smile, we need at least two factors.
  - In the two-factor case, there are 8 parameters.
- When calibrating, we find that the two-factor Bergomi model is already over-parameterized. Any combination of parameters that gives a roughly $1/\sqrt{T}$ ATM skew fits well enough.
  - Moreover, the calibrated correlations between the Brownian increments $dW_s^{(i)}$ tend to be high.

ATM skew in the Bergomi model

- 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 functional $C^{x\xi}$.
  - Which is in turn driven by the exponential kernel in the exponent in (3).
- The observed $\psi(\tau) \sim \tau^{-\alpha}$ for some α.
- It's tempting to replace the exponential kernels in (3) with a power-law kernel.

Tinkering with the Bergomi model

- 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.

The RFSV model

In [Gatheral, Jaisson and Rosenbaum], using RV estimates as proxies for daily spot volatilities, we uncovered two startlingly simple regularities:

- Distributions of increments of log volatility are close to Gaussian, consistent with many prior studies.
- For reasonable timescales of practical interest, the time series of volatility is consistent with the simple model

$$\log \sigma_{t+\Delta} - \log \sigma_t = \nu \left(W^H_{t+\Delta} - W^H_t \right) \tag{4}$$

  where W^H is fractional Brownian motion.
- We call our stationary version of (4) the Rough Fractional Stochastic Volatility (RFSV) model after the formally identical FSV model of [Comte and Renault].

Representation in terms of Brownian motion

The Mandelbrot-Van Ness representation of fractional Brownian motion W^H is as follows (with $\gamma = 1/2 = H$):

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

where the choice $C_H = \sqrt{\frac{2\,H\,\Gamma(3/2-H)}{\Gamma(H+1/2)\,\Gamma(2-2\,H)}}$ ensures that

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

Substituting into (4) (and in terms of $v_t = \sigma_t^2$), we obtain

$$\log v_u - \log v_t = 2\,\nu\,C_H \left\{ \int_{-\infty}^{u} \frac{dW_s^{\mathbb{P}}}{(u-s)^\gamma} - \int_{-\infty}^{t} \frac{dW_s^{\mathbb{P}}}{(t-s)^\gamma} \right\}.$$

Pricing under rough volatility (under $\mathbb{P}$)

Then

$$\begin{aligned} & \log v_u - \log v_t \\ = \quad & 2\,\nu\,C_H \left\{ \int_t^u \frac{1}{(u-s)^\gamma}\, dW_s^{\mathbb{P}} + \int_{-\infty}^t \left[\frac{1}{(u-s)^\gamma} - \frac{1}{(t-s)^\gamma} \right] dW_s^{\mathbb{P}} \right\} \\ =: \quad & 2\,\nu\,C_H \left\{ M_t(u) + Z_t(u) \right\}. \end{aligned} \tag{5}$$

- 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 forward variance curve

Taking the exponential of (5) gives

$$v_u = v_t \exp\left\{2\,\nu\,C_H\,\left[M_t(u) + Z_t(u)\right]\right\} \tag{6}$$

Ignoring the difference between $\mathbb{P}$ and $\mathbb{Q}$, and computing the conditional expectation gives

$$\begin{aligned}\mathbb{E}^{\mathbb{P}}\left[v_u \middle| \mathcal{F}_t\right] &= \xi_t(u) \\ &= v_t\,\exp\left\{2\,\nu\,C_H\,Z_t(u)\right\}\,\mathbb{E}\left[\exp\left\{2\,\nu\,C_H\,M_t(u)\right\}\middle|\,\mathcal{F}_t\right]\end{aligned}$$

where (by definition) $\xi_t(u)$ is the forward variance curve at time t.

- The $Z_t(u)$ are encoded in the forward variance curve $\xi_t(u)$!

The rough Bergomi model

Rewriting (6) gives

$$\begin{aligned} v_u &= \xi_t(u) \frac{\exp\{2\,\nu\, C_H\, M_t(u)\}}{\mathbb{E}\left[\exp\{2\,\nu\, C_H\, M_t(u)\}\middle|\, \mathcal{F}_t\right]} \\ &= \xi_t(u)\, \mathcal{E}\left(\eta\, \tilde{W}_t^H(u)\right) \end{aligned} \tag{7}$$

where $\eta = 2\,\nu\, C_H/\sqrt{2\,H}$, $\mathcal{E}(\cdot)$ denotes the stochastic exponential and

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

is a Volterra process with the fBm-like scaling property

$$\mathrm{var}\left[\tilde{W}_t^H(u)\right] = (u-t)^{2\,H}.$$

- We could call the model (7) a *rough Bergomi* or *rBergomi* model.

Features of the rough Bergomi model

- The forward variance curve

$$\xi_u(t) = \mathbb{E}\left[v_u \middle| \mathcal{F}_t \right] = v_t \, \exp\left\{ \eta \sqrt{2\,H} \, Z_t(u) + \frac{1}{2} \eta^2 \, (u - t)^{2\,H} \right\}.$$

  depends on the historical path $\{W_s, s < t\}$ of the Brownian motion since inception ($s = -\infty$ say).
- The rough Bergomi model is non-Markovian:

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

- However, given the (infinite) state vector $\xi_t(u)$, which can in principle be computed from option prices, the dynamics of the model are well-determined.

Re-interpretation of the conventional Bergomi model

- A conventional n-factor Bergomi model is not self-consistent for an arbitrary choice of the initial forward variance curve $\xi_t(u)$.
  - $\xi_t(u) = \mathbb{E}\left[v_u \middle| \mathcal{F}_t \right]$ should be consistent with the assumed dynamics.
- Viewed from the perspective of the fractional Bergomi model however:
  - The initial curve $\xi_t(u)$ reflects the history $\{W_s; s < t\}$ of the driving Brownian motion up to time t.
  - The exponential kernels in the exponent of (3) approximate more realistic power-law kernels.
- The conventional two-factor Bergomi model is then justified in practice as a tractable Markovian engineering approximation to a more realistic rough Bergomi model.

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.
- Then

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

with

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

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

The rBergomi model: Full specification

To recap, the rBergomi model reads:

$$\begin{aligned} \frac{dS_t}{S_t} &= \sqrt{v_t}\, dZ_t \\ v_u &= \xi_t(u)\, \mathcal{E}\left(\eta\, \tilde{W}_t^H(u)\right) \end{aligned} \tag{8}$$

with

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

- Note in particular that we have achieved our earlier wish to replace the exponential kernels in the Bergomi model with a power-law kernel.

Simulation of the rBergomi model

We simulate the rBergomi model as follows:

- For each time, construct the joint covariance matrix for the Volterra process $\tilde{W}^H$ and the Brownian motion Z.
- Generate iid normal random variables and perform a Cholesky decomposition to get a matrix of paths of $\tilde{W}^H$ and Z with the correct joint marginals.
- Hold these paths in memory to generate option prices for each expiration.
- Compute implied volatilities to get the volatility surface.
- This procedure is very slow!
    - Speeding up the simulation is work in progress.

Calibration of the rBergomi model

- The rBergomi model has only three parameters: H, η and ρ.
- If we had a fast simulation, we could just iterate on these parameters to find the best fit to observed option prices. But we don't.
- Alternatively, we could estimate H either from the term structure of ATM SPX skew, or from the term structure of *ATM* VIX volatilities.
  - Implied volatility of VIX should be "volatility of SPX volatility"!
- As we will see, even without proper calibration (*i.e.* just guessing parameters), rBergomi model fits to the volatility surface are amazingly good.

SPX smiles in the rBergomi model

- In Figure 9, we show how a rBergomi model simulation is consistent with the SPX option market as of 04-Feb-2010, a day when the ATM volatility term structure happened to be pretty flat.
  - rBergomi parameters were: $H = 0.07$, $\eta = 1.9$, $\rho = -0.9$.
  - Only three parameters to get a very good fit to the whole SPX volatility surface!

rBergomi fits to SPX smiles as of 04-Feb-2010

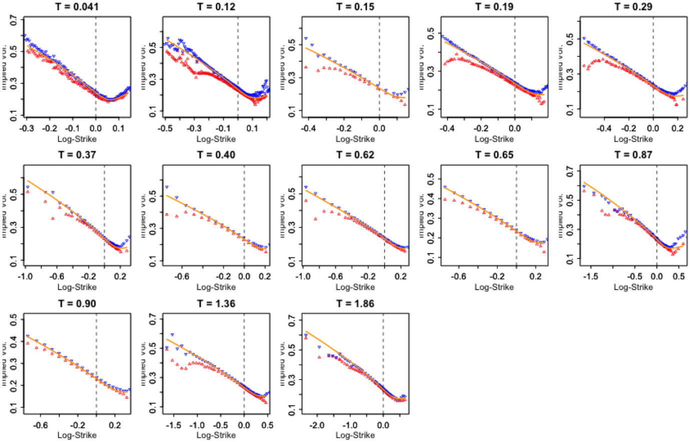

Figure 8: Red and blue points represent bid and offer SPX implied volatilities; orange smiles are from the rBergomi simulation.

rBergomi 04-Feb-2010 fit detail

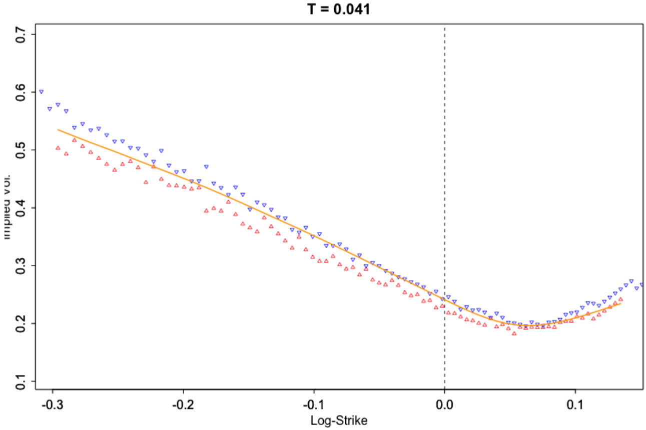

Figure 9: Feb-2014 expiration: Red and blue points represent bid and offer SPX implied volatilities; orange smile is from the rBergomi simulation.

rBergomi ATM skews and volatilities

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

Term structure of ATM skew in the rBergomi model

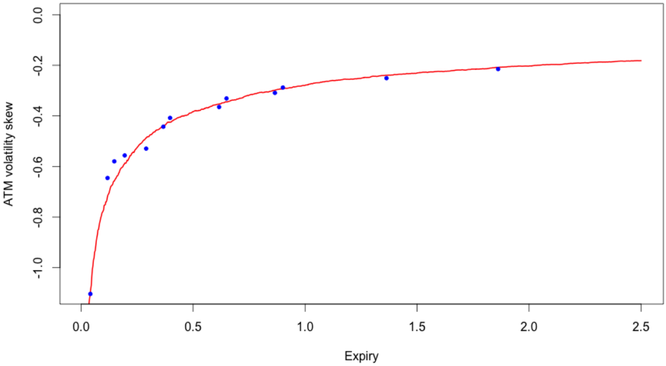

Figure 10: Blue points are empirical skews; the red line is from the rBergomi simulation.

Term structure of ATM volatility in the rBergomi model

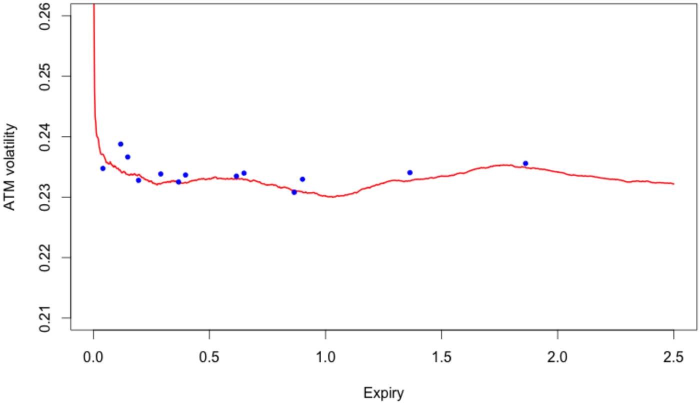

Figure 11: Blue points are empirical ATM volatilities; the red line is from the rBergomi simulation.

VIX smiles as of 04-Feb-2010

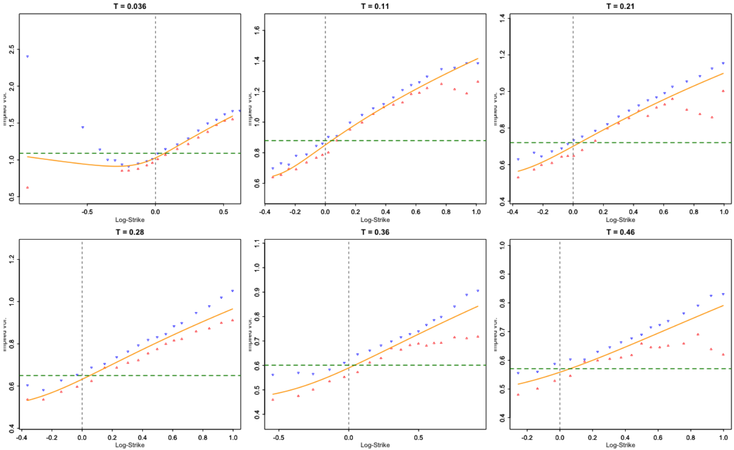

Figure 12: Blue ask volatilities; red points are bid volatilities; orange lines are SVI fits; green dashed lines represent the VIX log-strip (VVIX).

VIX futures in the rBergomi model

We denote the square of the VIX futures payoff by

$$\psi(T) = \frac{1}{\Delta} \int_T^{T+\Delta} \mathbb{E}[v_u|\mathcal{F}_T]\, du.$$

Then

$$\mathbb{E}\left[\psi(T)|\, \mathcal{F}_t\right] = \frac{1}{\Delta} \int_T^{T+\Delta} \xi_t(u)\, du.$$

and, approximating the arithmetic mean by the geometric mean,

$$\mathrm{var}[\log \psi(T)|\mathcal{F}_t] \approx \eta^2\, D_H^2\, (T-t)^{2\,H}\, f^H\left(\frac{\Delta}{T-t}\right)$$

where $D_H = \frac{\sqrt{2\,H}}{H+1/2}$ and

$$f^H(\theta) = \frac{1}{\theta^2} \int_0^1 \left[(1+\theta-x)^{1/2+H} - (1-x)^{1/2+H}\right]^2 dx.$$

VVIX term structure in the rBergomi model

The VIX variance swaps ($VVIX^2$) are then given by

$$\begin{aligned} VVIX^2(T)\,(T-t) &\approx \operatorname{var}\left[\log\sqrt{\psi(T)}\,\middle|\,\mathcal{F}_t\right] \\ &\approx \frac{1}{4}\,\eta^2\,D_H^2\,(T-t)^{2\,H}\,f^H\left(\frac{\Delta}{T-t}\right). \quad (9) \end{aligned}$$

VVIX term structure as of 04-Feb-2010

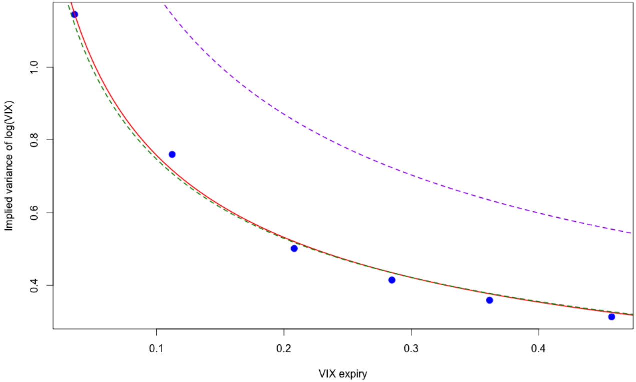

Figure 13: Blue points are implied QV of $\log(VIX)$ ($VVIX^2$). The red curve is a fit of equation (9) with $\eta = 2.36$ and $H = 0.022$; the purple dashed curve corresponds to $H = 0.07$ and $\eta = 1.9$; the green dashed curve corresponds to $H = 0.036$ and $\eta = 1.9$.

Remarks on the VVIX fit

- Recall that the parameters we used to get the SPX fit were just guessed.
  - The VIX fit indicates that we should perhaps decrease H, increase η and decrease ρ (to maintain the original skew levels).
  - Although we also see that it is hard to determine H and η individually.
- Recall also that the rBergomi model generates flat VIX smiles.
  - So although rBergomi may fit the term structure of VVIX, it is not consistent with observed VIX smiles.

Change of measure

- So far, we have conflated $\mathbb{P}$ and $\mathbb{Q}$.
- We know, in particular from VIX options, that these two measures are different.
  - Like the conventional Bergomi model, the rBergomi model predicts VIX smiles that are almost exactly flat.
  - In contrast, observed VIX smiles are strongly upward sloping (see Figure 13).
- Intuitively, high volatility scenarios are priced more highly by the market than low volatility scenarios.

Formulation of the rough volatility model under $\mathbb{Q}$

With some general change of measure

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

we may rewrite the rough Bergomi model (7) under $\mathbb{Q}$ as

$$\begin{aligned} v_u &= \mathbb{E}^{\mathbb{P}}\left[v_u \middle| \mathcal{F}_t\right] \exp\left\{\eta\sqrt{2H}\int_t^u \frac{1}{(u-s)^\gamma}\, dW_s^{\mathbb{P}} - \eta^2\,(u-t)^{2H}\right\} \\ &= \mathbb{E}^{\mathbb{P}}\left[v_u \middle| \mathcal{F}_t\right] \exp\left\{\eta\sqrt{2H}\int_t^u \frac{1}{(u-s)^\gamma}\, dW_s^{\mathbb{Q}} - \eta^2\,(u-t)^{2H}\right. \\ &\qquad \left. +\,\eta\sqrt{2H}\int_t^u \frac{1}{(u-s)^\gamma}\,\mu_s\, ds\right\}. \end{aligned}$$

The last term in the exponent obviously changes the marginal distribution of the v_u; the v_u will not be lognormal in general under $\mathbb{Q}$.

The simplest change of measure

With the simplest change of measure

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

with μ constant, we would have (still rBergomi form)

$$\begin{aligned} v_u &= \mathbb{E}^{\mathbb{P}}\left[v_u \middle| \mathcal{F}_t \right] \mathcal{E}\left(\eta \sqrt{2\,H} \int_t^u \frac{dW_s^{\mathbb{Q}}}{(u-s)^{\gamma}} \right) \\ &\qquad \times \exp\left\{ \eta \sqrt{2\,H} \int_t^u \frac{\mu}{(u-s)^{\gamma}}\, ds \right\} \\ &= \xi_t(u)\, \mathcal{E}\left(\eta \sqrt{2\,H} \int_t^u \frac{dW_s^{\mathbb{Q}}}{(u-s)^{\gamma}} \right) \end{aligned} \tag{10}$$

where

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

Estimating the price of volatility risk

Performing the last integration explicitly gives

$$\xi_t(u) = \mathbb{E}^{\mathbb{P}}\left[v_u \middle| \mathcal{F}_t \right] \exp\left\{ \mu\, \eta \, \frac{\sqrt{2\,H}}{H+1/2}\,(u-t)^{H+1/2} \right\}. \qquad (11)$$

- $\mathbb{E}^{\mathbb{P}}\left[v_u \middle| \mathcal{F}_t \right]$ is a variance forecast computed from the history of RV as explained in [Gatheral, Jaisson and Rosenbaum].
- $\xi_t(u)$ may be estimated using the log-strip of SPX options.
- In principle, we could compare the two historically to estimate the price of risk μ using (11).

Empirical study

For each of 2,658 days from Jan 27, 2003 to August 31, 2013:

- We compute proxy variance swaps from closing prices of SPX options sourced from OptionMetrics (www.optionmetrics.com) via WRDS.
- We form the forecasts $\mathbb{E}^{\mathbb{P}}\left[v_u \middle| \mathcal{F}_t\right]$ using 500 lags of SPX RV data sourced from The Oxford-Man Institute of Quantitative Finance (http://realized.oxford-man.ox.ac.uk).
    - To adjust for overnight variance (RV is over the trading day only), we could follow [Corsi, Fusari, and La Vecchia] by rescaling these estimates to match the sample close-to-close variance.
        - This scaling factor would then be 1.3658.
    - However, we quickly observe that η, H and the scaling factor are all time-varying quantities.

The RV scaling factor

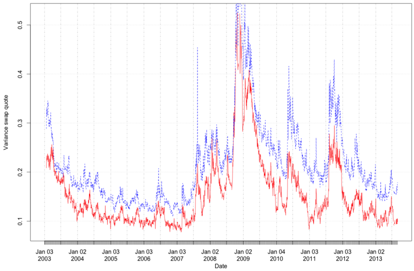

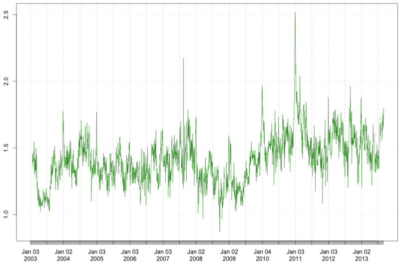

Figure 14: The LH plot shows actual (proxy) 3-month variance swap quotes in blue vs forecast in red (with no scaling factor). The RH plot shows the ratio between 3-month actual variance swap quotes and 3-month forecasts.

More on scaling

- Empirically, it seems that the variance curve is a simple scaling factor times the forecast, but that this scaling factor is time-varying.
- Recall that as of the close on Friday September 12, 2008, it was widely believed that Lehman Brothers would be rescued over the weekend. By Monday morning, we knew that Lehman had failed.
- In Figure 15, we see that variance swap curves just before and just after the collapse of Lehman are just rescaled versions of the RFSV forecast curves.

Actual vs predicted over the Lehman weekend

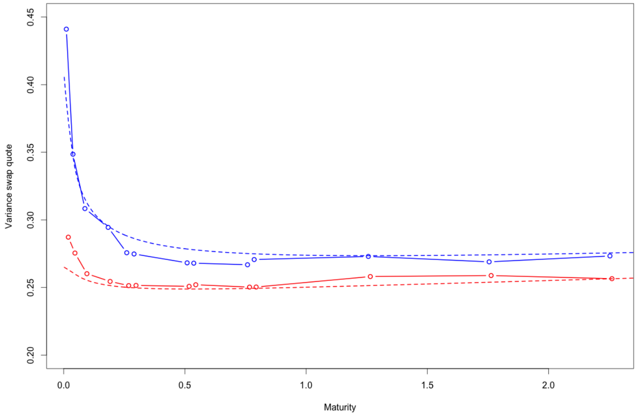

Figure 15: S&P variance swap curves as of September 12, 2008 (red) and September 15, 2008 (blue). The dashed curves are RFSV model forecasts rescaled by the average 3-month scaling factor over the prior week.

Remarks

We note that

- The actual variance swaps curves are very close to the forecast curves, up to a scaling factor.
- We are able to explain the change in the variance swap curve with only one extra observation: daily variance over the trading day on Monday 15-Sep-2008.
- The SPX options market appears to be backward-looking in a very sophisticated way.

Summary

- We uncovered a remarkable monofractal scaling relationship in historical volatility.
- This leads to a natural non-Markovian stochastic volatility model under $\mathbb{P}$.
- The simplest specification of $\frac{d\mathbb{Q}}{d\mathbb{P}}$ 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 only three parameters.
- For perhaps the first time, we have a simple consistent model of historical and implied volatility.

References

Lorenzo Bergomi, Smile dynamics II, *Risk Magazine* 67–73 (October 2005).

Lorenzo Bergomi and Julien Guyon, The smile in stochastic volatility models, *SSRN* (2011).

Fabienne Comte and Eric Renault, Long memory in continuous-time stochastic volatility models, *Mathematical Finance* **8** 29–323(1998).

F. Corsi, N. Fusari, and D. La Vecchia, Realizing smiles: Options pricing with realized volatility, *Journal of Financial Economics* **107**(2) 284–304 (2013).

Jim Gatheral, *The Volatility Surface: A Practitioner's Guide*, Wiley Finance (2006).

Jim Gatheral and Antoine Jacquier, Arbitrage-free SVI volatility surfaces, *Quantitative Finance*, **14**(1) 59–71 (2014).

Jim Gatheral, Thibault Jaisson and Mathieu Rosenbaum, Volatility is rough, *Available at SSRN 2509457*, (2014).