

Rough volatility: evidence from option prices

Giulia Livieri
Scuola Normale Superiore, Pisa
giulia.livieri@sns.it

Saad Mouti
University Pierre and Marie Curie (Paris 6)
saad.mouti@upmc.fr

Andrea Pallavicini
Imperial College London and Banca IMI
andrea.pallavicini@imperial.ac.uk

Mathieu Rosenbaum
École Polytechnique
mathieu.rosenbaum@polytechnique.edu

9 February 2017

Abstract

It has been recently shown that spot volatilities can be very well modeled by rough stochastic volatility type dynamics. In such models, the log-volatility follows a fractional Brownian motion with Hurst parameter smaller than $1/2$. This result has been established using high frequency volatility estimations from historical price data. We revisit this finding by studying implied volatility based approximations of the spot volatility. Using at-the-money options on the S&P500 index with short maturity, we are able to confirm that volatility is rough. The Hurst parameter found here, of order 0.3, is slightly larger than that usually obtained from historical data. This is easily explained from a smoothing effect due to the remaining time to maturity of the considered options.

Keywords: Rough volatility, fractional Brownian motion, implied volatility, Medvedev-Scaillet approximation.

1 Introduction

Since the seminal work of Black and Scholes [4], the most classical way to model the behavior of the price S_t of a financial asset is to use a continuous

semi-martingale dynamic of the form

$$d \log S_t = \mu_t dt + \sigma_t dW_t,$$

with μ_t a drift process and W_t a Brownian motion. The coefficient σ_t is referred to as the volatility process. As is well-known, it is the key ingredient in the model when one is interested in derivatives pricing and hedging.

Historically, following the pioneering approach of [4], practitioners have first considered the case where the process σ_t is constant or deterministic, that is the Black and Scholes model. However, in the late eighties, it became clear that such specification for the volatility is inadequate. In particular, the Black and Scholes model is inconsistent with observed prices for liquid European options. Indeed the implied volatility, that is the volatility parameter that should be plugged into the Black-Scholes formula to retrieve a market option price, depends in practice on the strike and maturity of the considered option, whereas it is constant in the Black-Scholes framework.

Hence more sophisticated models have been introduced. A first possible extension, proposed by Dupire [7] and Derman and Kani [6], is to take σ_t as a deterministic function of time and asset price. Such models, called local volatility models, enable us to perfectly reproduce a given implied volatility surface. However, its dynamic is usually quite unrealistic under local volatility. Another approach is to consider the volatility σ_t itself as an Ito process driven by an additional Brownian motion, typically correlated to W . Doing so one obtains less accurate static fits for the implied volatility surface but more suitable dynamics. Among the most famous of these stochastic volatility models are the Hull and White model [17], the Heston model [16] and the SABR model [15]. More recent market practice is to use so-called local-stochastic volatility models which both fit the market exactly and generate reasonable dynamics.

In all the Brownian volatility models mentioned above, the smoothness of the sample path of the volatility is the same as that of a Brownian motion, namely $1/2 - \varepsilon$ Hölder continuous, for any $\varepsilon > 0$. However, it is shown in [13] that in practice, spot volatility is much rougher than this. This result in [13] is based on a statistical analysis of historical data using sophisticated high frequency estimation methods. More precisely, it is established in [13] that the dynamic of the log-volatility process is very close to that of a fractional Brownian motion with Hurst parameter smaller than $1/2$. Recall that a fractional Brownian motion W^H with Hurst parameter $H \in (0, 1)$ is a Gaussian process with stationary increments such that

$$\text{Cov}[W_t^H, W_s^H] = \frac{1}{2} (|t|^{2H} + |s|^{2H} - |t - s|^{2H}).$$

The Hölder regularity of W^H is $H - \varepsilon$ for any $\varepsilon > 0$ and for $H = 1/2$ we retrieve the classical Brownian motion. Therefore models where the volatility is driven by a fractional Brownian motion with $H < 1/2$ are called rough volatility models. Beyond fitting almost perfectly historical volatility time series, rough volatility models enable us to reproduce important stylized facts of liquid option prices that local/stochastic volatility models typically fail to generate. In particular, the exploding term structure when maturity goes to zero of the at-the-money skew (the derivative of the implied volatility with respect to strike) is readily obtained, see [1, 11]. Other developments about rough volatility models can be found in [2, 3, 8, 9, 10, 12, 14, 18, 21].

The goal of this paper is to revisit the finding in [13] using implied volatility data. Indeed in [13], the authors work with historical price data from underlyings to estimate spot volatility. Here we use a spot volatility proxy which is not based on historical data, but on implied volatility. More precisely, we approximate the spot volatility by the implied volatility of an at-the-money liquid option with short maturity (or a refined version of it). This idea can be justified by the fact that in most models, the at-the-money implied volatility tends to the spot volatility as maturity goes to zero, see for example [20]. Our main result is a confirmation of that in [13]: When using alternate spot volatility measurement methods based on option prices, we can still conclude that volatility is rough.

The paper is organized as follows. We investigate in Section 2 the roughness of time series of spot volatility approximations given by implied volatilities of at-the-money options on the S&P500 index, with maturity one month. In Section 3, instead of using raw implied volatilities, we compute spot volatilities from implied ones through a correction formula due to Medvedev and Scaillet, see [19]. We then carry the same analysis as in Section 2. The results in Sections 2 and 3 are very similar to those in [13]. However, the estimated values for the Hurst parameter, although smaller than $1/2$, are actually larger than those obtained in [13]. We show numerically and analytically in Section 4 that this upward bias comes from a regularizing effect due to the remaining time to maturity of the considered options.

2 At-the-money implied volatility with short maturity as spot volatility proxy

As explained in the introduction, our goal is to study the behavior of the spot volatility and to show that it is well approximated by a rough process. Of course this is a difficult task since volatility is a latent, unobserved variable. In [13], the authors use recent estimation methods based on ultra high frequency price data to estimate spot volatility. In this work, instead of

using historical data as in [13], we wish to use option price data. This idea is reasonable if we use at-the-money options for which the time to maturity is short. Indeed, it is well-known that in most models, the at-the-money implied volatility converges to the spot volatility as maturity goes to zero, see for example [20].

2.1 Data description

In this section, we use a data set from Bloomberg¹, made of daily observations of the implied volatility of the option with maturity one month on the S&P500 index, from January 5, 2006 to May 5, 2011². Note that the data are in fact already extrapolated internally by the data provider (using quoted options at 4 PM) and do not necessarily exactly correspond to transaction data, see [5]. In Section 3, we present a method enabling us to derive spot volatilities from observed option prices with various maturities. Here we rely on the data provider approach to get option prices with the same maturity. This is not an issue since our aim in this work is to show that a rough dynamic for the volatility is obtained from any reasonable spot volatility proxy.

2.2 Scaling property

Let $\sigma_{t_0}^{imp}, \dots, \sigma_{t_N}^{imp}$ be the time series of implied volatilities extracted from our data base. Here for $i \geq 0$, $t_{i+1} - t_i$ corresponds to one business day. In the spirit of [13], we wish to review the behavior of the so-called structure function $m(q, \Delta)$ given by

$$m(q, \Delta) = \frac{1}{N} \sum_{k=0}^{\lfloor (N-1)/\Delta \rfloor} |\log(\sigma_{t_{(k+1)\Delta}}^{imp}) - \log(\sigma_{t_{k\Delta}}^{imp})|^q$$

for various $q > 0$ and lags Δ going from 1 to about 40 days³. Through the quantity $m(q, \Delta)$, our goal is to revisit the finding in [13] that the (spot) log-volatility is well approximated by a fractional Brownian motion with Hurst parameter H smaller than $1/2$. In this case, assuming spot and implied volatilities coincide, we should observe the following relationship:

$$m(q, \Delta) \sim c_q \Delta^{qH}, \tag{1}$$

with c_q a constant depending on q . Indeed, we have for $t \geq 0$ and $\Delta > 0$

$$\mathbb{E}[|W_{t+\Delta}^H - W_t^H|^q] = \tilde{c}_q \Delta^{qH},$$

¹Data obtained from AXA Group Risk Management.

²Data around the third Friday of each month (settlement date) are removed from the data base. We have 1166 points in total.

³Of course when computing $m(q, \Delta)$ we in fact also average over the possible starting points $t_0, \dots, t_{\Delta-1}$.

with \tilde{c}_q the absolute moment of order q of a standard Gaussian random variable.

To investigate the validity of (1), we plot in Figure 1 the logarithm of $m(q, \Delta)$ against the logarithm of Δ , for several values of q .

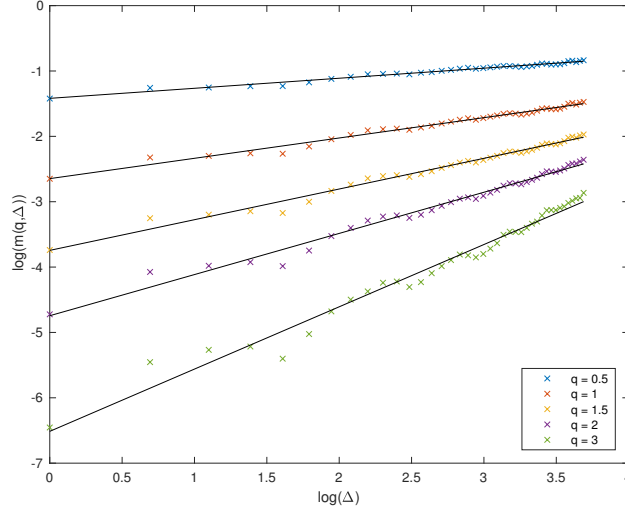


Figure 1: Scaling property of log-volatility increments.

For every q , the points with coordinates $(\log(\Delta), \log(m(q, \Delta)))$ are almost perfectly on the same line, and this for a wide range of Δ . Figure 1 is actually very similar to that obtained from historical volatility measurements in [13]. Thus we can deduce that indeed, for a given q ,

$$m(q, \Delta) \sim c_q \Delta^{\zeta(q)},$$

for some $\zeta(q)$.

Now we want to check whether $\zeta(q)$ can be taken of the form qH for some H , as suggested in [13]. This would lead to the same monofractal scaling as that of the fractional Brownian motion with Hurst parameter H . To answer this, we plot in Figure 2 the points with coordinates $(q, \zeta(q))$, where $\zeta(q)$ is taken as the slope of the line in Figure 1 corresponding to the power q , and the points with coordinates $(q, 0.32q)$.

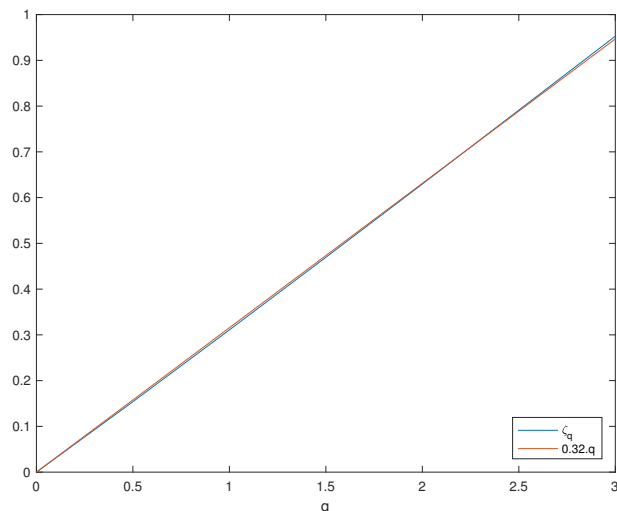


Figure 2: Monofractal scaling.

We see that the two graphs on Figure 2 can hardly be distinguished. This means that (1) almost perfectly holds, with H around 0.32. Note that such value for H corresponds to rough volatility since it is smaller than $1/2$. However, it is larger than those reported in [13]. This is actually due to the fact that our options have a significant remaining time to maturity of one month. This induces a smoothing phenomenon in the estimation of the Hurst parameter. This effect is of the same nature as that described and explained in [13] caused by the discrepancy between spot and integrated volatility over a short time interval. We quantify this measurement bias numerically and analytically in Section 4.

2.3 Distribution of log-volatility increments

Recall that it is suggested in [13] that the log-volatility process is well modeled by a fractional Brownian motion with Hurst parameter smaller than $1/2$. This implies monofractal scaling as investigated above but also a Gaussian behavior of the log-volatility increments. This feature is indeed satisfied when using historical estimates as measurements for spot volatility, see [13]. Here we wish to study whether such property also holds when the volatility proxies are given by our short term at-the-money implied volatilities. To this end, we display in Figure 3 histograms of log-volatility increments over different time intervals, together with a Gaussian density fit and the Gaussian density associated to the increments of a fractional Brownian motion with Hurst parameter equal to 0.32.

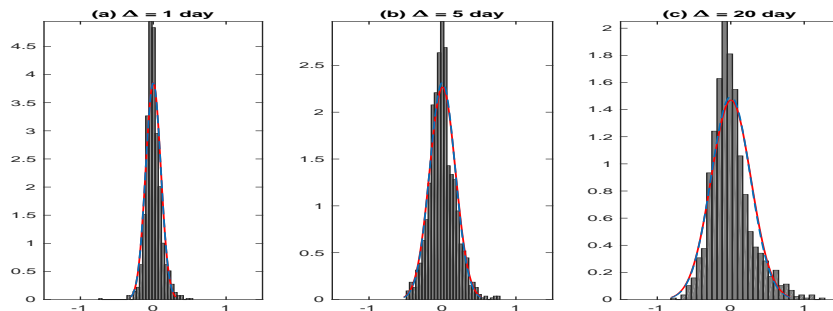


Figure 3: Distribution of the log-volatility increments when using implied volatility as spot volatility proxy. The Gaussian fit is in blue and the density associated to the increments of a fractional Brownian motion with Hurst parameter equal to 0.32 is in red.

From these graphs, we obtain that empirical distributions of log-volatility increments are reasonably approximated by Gaussian laws. However, we can remark that the empirical distributions are slightly more concentrated around their center. Finally, the Gaussian fits almost exactly coincide with those associated to the fractional Brownian motion with Hurst parameter equal to 0.32.

In conclusion, using at-the-money implied volatilities with maturity one month as spot volatility proxies, we obtain that log-volatility is well approximated by a rough fractional Brownian motion. This confirms the finding in [13].

3 A refined implied volatility based proxy for the spot volatility

In this section, we wish to study the robustness of the results obtained in Section 2. To do so, we work with another spot volatility proxy based on at-the-money options with short maturity. More precisely, we use the approximation formula from Medvedev and Scaillet, see [19]. This correction formula enables us to compute a spot volatility proxy from an at-the-money implied volatility with any (short) maturity. This is an advantage compared to what is done in Section 2 where only options with one month maturity are considered⁴. The drawback of Medvedev-Scaillet formula is that it is proved to be valid only within a restricted class of stochastic volatility models, which does not include rough volatility models. However our goal here is to see whether a proxy obtained from a Brownian volatility model still exhibits a rough behavior.

⁴Mixing various maturities without any correction would have been very arguable.

3.1 Data description and processing

Here our data set is provided by OptionMetrics and consists in daily close bid/ask prices of European puts and calls on the S&P500 index, from September 5, 2001 to January 31, 2012, for various strikes and maturities, together with the daily traded volumes. We discard options with price less than 2.5 cents of dollars or with zero trading volume. Besides, as in Section 2, prices corresponding to settlement dates are removed, so as obvious outliers.

We then want to compute implied volatilities from put and call prices. Thus we have to invert (everyday) the Black-Scholes formula. Therefore we need to fix for any time to maturity τ an underlying forward price $F(\tau)$ and a zero coupon bond price $D(\tau)$. To do so, we use the following classical approach based on put-call parity. The values of $F(\tau)$ and $D(\tau)$ are taken as solutions of the minimization problem

$$\arg \min_{D,F} \left\{ \sum_i w_i \left(\frac{1}{2}(C_i^a - P_i^b) + \frac{1}{2}(C_i^b - P_i^a) - D(\tau)(F(\tau) - K_i) \right) \right\},$$

where $C_i^{a,b}$ and $P_i^{a,b}$ are respectively the call and put market prices (a standing for ask, b for bid) quoted at strike level K_i . The weights w_i are given by

$$w_i = \frac{\sqrt{\min\{V_i^C, V_i^P\}}}{\frac{1}{2}(C_i^a - C_i^b) + \frac{1}{2}(P_i^a - P_i^b)},$$

with V_i^C and V_i^P the trading volumes of call and put options at strike K_i . Finally, our implied volatility is taken as that of a call whose price would be the midprice between the bid and ask prices.

Recall that for our approximations to be valid, we focus on at-the-money implied volatilities with short maturity. Following [19], we only select implied volatilities of options with time to maturity ranging from 15 to 60 days. Shorter term options are discarded because quotes can be noisy. Moreover, we restrict our data to log forward moneyness belonging to the interval $[-0.03, 0.03]$. Such procedure yields a total number of 34842 implied volatilities over 2569 days.

3.2 The Medvedev-Scaillet correction formula

In [19], the authors consider a general modeling framework encompassing most of the classical parametric price models. They use a two factors jump-diffusion stochastic volatility model of the form

$$\begin{cases} dS_t = (r - \mu(\sigma_t)) S_t dt + \sigma_t S_t dZ_t + S_t dJ_t \\ d\sigma_t = a(\sigma_t) dt + b(\sigma_t)(\rho dZ_t + \sqrt{1 - \rho^2} dW_t), \end{cases} \quad (2)$$

where Z_t and W_t are two independent Brownian motions and J_t is a Poisson-type jump process, independent of Z_t and W_t . Both r and the correlation coefficient ρ are assumed to be constant. The expected jump size $\mathbb{E}[\Delta J]$ is also constant, but the jump intensity $\lambda(\sigma_t)$ may depend on the volatility in a deterministic way. Here, as in the numerical experiments in [19], we consider the following parametric forms:

$$b(\sigma_t) = \beta\sigma_t^\phi, \quad \lambda(\sigma_t) = \lambda_0\sigma_t^\psi,$$

for some non-negative constants β , ϕ , λ_0 and ψ .

Let σ be the spot volatility and $\hat{\sigma} = \hat{\sigma}(\tau)$ be the at-the-money implied volatility of an option with time to maturity τ . Following [19], we build up our option-based spot volatility proxy in two steps. First, the chosen model is calibrated from the approximation formula in Proposition 7 in [19] using all our option prices over the entire time period. To retrieve the proxy for the spot volatility, we then consider the following expansion as τ goes to zero shown in [19]:

$$\begin{aligned} \sigma &= \hat{\sigma} - I_1(0, \hat{\sigma})\sqrt{\tau} \\ &+ \left(I_1(0, \hat{\sigma})\frac{\partial I_1(0, \hat{\sigma})}{\partial \sigma} - I_2(0, \hat{\sigma}) + \frac{1}{2}\rho b(\hat{\sigma})\mathbb{E}[\Delta J]\frac{\partial \lambda(\hat{\sigma})}{\partial \sigma} \right)\tau + O(\tau\sqrt{\tau}). \end{aligned} \quad (3)$$

The functions I_1 and I_2 are explicitly defined in [19] and depend only on β , ρ , ϕ , λ_0 , ψ and $\mathbb{E}[\Delta J]$.

3.3 The scaling property revisited

We now wish to study the scaling property of spot volatility proxies based on the approximation formula (3). We consider two cases: The Heston case, where $\phi = 0$ and $\lambda_0 = 0$, and the general case, where all the parameters are calibrated. The calibration results are given in Table 1.

PARAMETER	HESTON	GENERAL CASE
$\beta\rho$	-0.18 (0.00)	-3.27 (0.08)
ρ	-0.48 (0.00)	-0.39 (0.00)
ϕ	0	1.79 (0.02)
$\lambda_0\mathbb{E}(\Delta J)$	0	-0.6924 (0.03)
$\mathbb{E}(\Delta J)$	---	-0.17 (0.00)
ψ	---	1.11 (0.01)

Table 1: Parameters calibrated on quoted S&P500 option prices, from September 5, 2001 to January 31, 2012.

Once the parameters are obtained, we can implement Equation (3) to compute everyday a spot volatility proxy. Note that in Equation (3), we take

for $\hat{\sigma}$ the implied volatility with shortest time to maturity. Then we conduct the same analysis as in Section 2.2. The results are given in Figure 4 for the Heston model and Figure 5 for the general case (notations are the same as in Section 2.2).

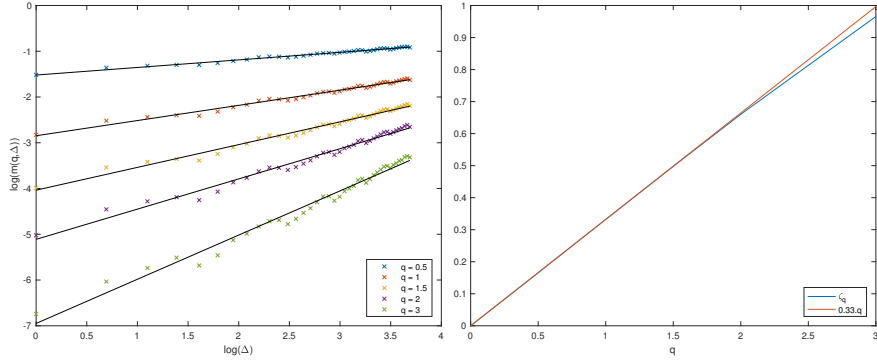


Figure 4: Scaling property of log-volatility increments when based on Heston proxy. In the second graph H is taken equal to 0.33.

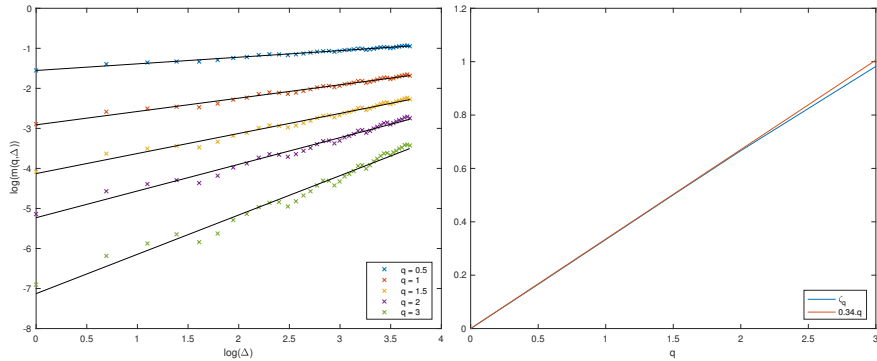


Figure 5: Scaling property of log-volatility increments when based on the general case proxy. In the second graph H is taken equal to 0.34.

The results are very similar to those in Section 2.2. Here again we can confirm the fact that volatility is rough. This is even obtained although in the models in which the proxies are computed, volatility is of Brownian type and therefore not rough.

4 On the upward bias when estimating the Hurst parameter

We explain in this section why using implied volatility measures as spot volatility proxies induces an upward bias in the estimation of the Hurst parameter. We start with a numerical investigation of this phenomenon.

4.1 Monte Carlo study

To understand the extend of the bias when estimating the Hurst parameter, we simulate option prices in a rough volatility model. Then we compute the Hurst parameter based on these simulated data. Let $T > 0$. We consider the following model without leverage effect over the time interval $[0, T]$:

$$d \log S_t = \sigma_t dZ_t, \quad d \log \sigma_t = \eta dW_t^H.$$

Here Z_t is a Brownian motion, W_t^H a fractional Brownian motion independent of Z_t and $\eta > 0$.

4.1.1 Simulation of fractional Brownian motion

We consider a time interval $[0, T]$ and fix an equidistant partition $0 = t_0 < t_1 < \dots < t_n = T$. We first wish to simulate $(W_{t_1}^H, \dots, W_{t_n}^H)$. For $i, j \in \{1, \dots, n\}$, we have

$$\mathbb{E}[W_{t_i}^H W_{t_j}^H] = \frac{1}{2} (t_i^{2H} + t_j^{2H} - |t_i - t_j|^{2H}).$$

Then we can use the Cholesky decomposition of the covariance matrix Σ of $(W_{t_1}^H, \dots, W_{t_n}^H)$: $\Sigma = LL^T$, where $L = (l_{ij})_{i,j \in \{1,n\}}$ is lower-triangular. Thus simulating a sample path of the fractional Brownian motion at times (t_i) can be done generating a vector $X = (X_1, \dots, X_n)$ of independent standard Gaussian random variables and setting $(W_{t_1}^H, \dots, W_{t_n}^H) = LX$.

4.1.2 Simulating option prices under rough volatility

We place ourselves at time $t_i > 0$ and assume past spot volatilities and prices have been observed at times t_1, \dots, t_i . We want to compute the price at time t_i of an option with expiration date $t_k = t_i + \tau$ for some $\tau > 0$. The procedure goes as follows:

- We generate M paths of the volatility process on the interval $[t_{i+1}, t_k]$. This is done simulating $(W_{t_j}^H)_{t_{i+1} \leq t_j \leq t_k}$ conditional on past information, that is the filtration generated by $(X_{t_1}, \dots, X_{t_i})$. Using the lower triangular form of L , these new values for the fractional Brownian motion at times $t_{i+1} \leq t_j \leq t_k$ can be obtained writing

$$W_{t_j}^H = \sum_{p=1}^i l_{jp} X_p + \sum_{p=i+1}^j l_{jp} X_p.$$

The i first variables X_p are those used to simulate the fractional Brownian motion up to time t_i , whereas (X_{i+1}, \dots, X_j) is a sample of independent standard Gaussian random variables, independent from past values. Taking the exponential, we get our spot volatility sample path. We write σ^m for the m -th volatility trajectory.

- The price at time t_i of an at-the-money option with time to maturity τ is obtained computing

$$\frac{1}{M} \sum_{m=1}^M C_{BS} \left(S_{t_i}, \tau, \sqrt{\frac{1}{\tau} \sum_{p=i+1}^k (\sigma_{t_p}^m)^2} \right),$$

where $C_{BS}(S_{t_i}, \tau, \sigma)$ is the price of an at-the-money option with time to maturity τ in a Black-Scholes model with volatility σ , zero interest rate, and underlying value S_{t_i} .

- Eventually we invert Black-Scholes formula to obtain the implied volatility.

4.1.3 Results

We consider the following set of parameters: $H = 0.04$, $\eta = 1.0$ and $T = 1000$ days. Such parameters are consistent with [1, 13]. We take $\tau \in \{1, \dots, 20\}$ days and run $M = 10^4$ simulations. Figure 6 displays the sample path of the spot volatility together with those of the implied volatilities associated to 5 and 20 days.

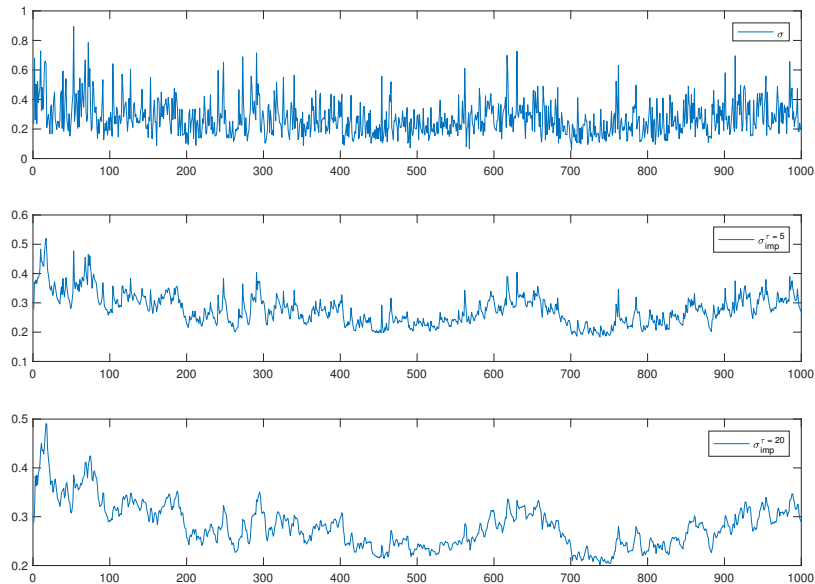


Figure 6: Sample paths of spot volatility and implied volatilities for $\tau = 5$ and $\tau = 20$.

At the visual level, it is already clear that implied volatility trajectories are not as rough as that of the spot volatility. Furthermore, the longer the time to maturity, the larger the smoothing effect.

As in Sections 2 and 3, we now consider Equation (1). Based on our simulation, for several values of q , we plot in Figure 7 the logarithm of $m(q, \Delta)$ against the logarithm of Δ . This is done in two cases: when m is obtained from spot volatility values and when m is derived from implied volatility values, with $\tau = 5$ days.

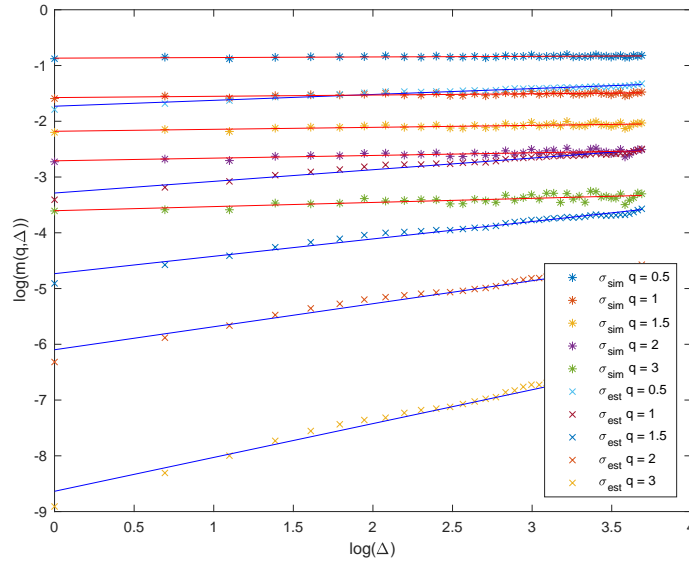


Figure 7: Scaling property of log-volatility increments: spot volatility and implied volatility with $\tau = 5$.

We see that for a given q , when $m(q, \Delta)$ is computed from implied volatilities, the points with coordinates $(\log(\Delta), \log(m(q, \Delta)))$ remain on the same line. However, the slope of this line is larger than that obtained when $m(q, \Delta)$ is computed from spot volatilities (which provides the true underlying H up to small statistical error). Hence there is indeed a smoothing effect due to the remaining time to maturity of the considered options.

Finally, we give in Figure 8 the estimated values of H when using implied volatilities from the simulation, for different times to maturity.

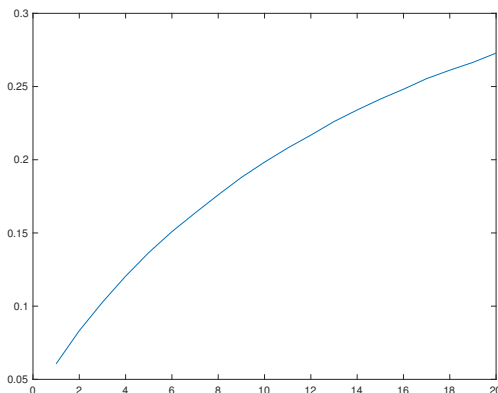


Figure 8: Estimated values of the Hurst parameter using implied volatilities as a function of time to maturity.

Under our simulation framework, we see that using options with maturity 1 day, we obtain a quite accurate value for H of 0.06, while the true parameter is equal to 0.04. Taking longer maturities leads to an increasing bias. With 20 days maturity, one gets an estimated Hurst parameter of about 0.27. These results are in line with those in Sections 2 and 3.

4.2 Analytical illustration of the upward bias

In the spirit of Appendix C in [13], we finally want to provide a more quantitative understanding of the observed upward bias when estimating the Hurst parameter from implied volatilities. To do so, we consider a very crude approximation. Indeed we suppose that the at-the-money implied variance at time t of an option with time to maturity $\tau > 0$, denoted by $\hat{v}^\tau(t)$, is given by

$$\hat{v}^\tau(t) = \frac{1}{\tau} \int_t^{t+\tau} \mathbb{E}_t[v_u] du,$$

where v_u is the spot variance at time u and $\mathbb{E}_t[\cdot]$ the conditional expectation operator with respect to information up to time t . Furthermore, we take a simplified rough volatility model assuming that for $u > 0$,

$$v_u = v_0 + \nu W_u^H,$$

for some $v_0 > 0$ and $\nu > 0$. These approximations are actually probably enough to shed light on the bias phenomenon. Indeed it is due to the effects of the conditional expectation and integral operators appearing in the implied volatility.

In this simplified setting, our goal is to illustrate the smoothing effect leading to the upward bias. To do so, we compute a quantity very related to $m(2, \Delta)$,

namely

$$\hat{m}^\tau(2, \Delta) = \mathbb{E}[(\hat{v}^\tau(\Delta) - \hat{v}^\tau(0))^2].$$

Indeed, under our assumptions, if the implied volatility were equal to the spot one, this quantity would be proportional to Δ^{2H} . However, we now show that because of the use of implied volatility in $\hat{m}(2, \Delta)$, this relationship no longer holds, particularly for large τ/Δ .

We recall the Mandelbrot and Van Ness representation of fractional Brownian motion:

$$W_t^H = c_H \left(\int_0^t (t-s)^{H-1/2} dW_s + \int_{-\infty}^0 ((t-s)^{H-1/2} - (-s)^{H-1/2}) dW_s \right),$$

where W_t is a two-sided Brownian motion and c_H is so that the variance of W_1^H is equal to 1. We easily have

$$\begin{aligned} \hat{v}^\tau(\Delta) &= v_0 + \frac{\nu}{\tau} c_H \int_0^\tau \int_{-\infty}^0 ((\Delta + u - s)^{H-1/2} - (-s)^{H-1/2}) dW_s du \\ &\quad + \frac{\nu}{\tau} c_H \int_0^\tau \int_0^\Delta (\Delta + u - s)^{H-1/2} dW_s du. \end{aligned}$$

Using stochastic Fubini theorem, this gives

$$\begin{aligned} \hat{v}^\tau(\Delta) - \hat{v}^\tau(0) &= \frac{\nu}{\tau} c_H \int_{-\infty}^0 \int_0^\tau ((\Delta + u - s)^{H-1/2} - (u - s)^{H-1/2}) du dW_s \\ &\quad + \frac{\nu}{\tau} c_H \int_0^\Delta \int_0^\tau (\Delta + u - s)^{H-1/2} du dW_s. \end{aligned}$$

Hence we easily deduce from Ito isometry that

$$\hat{m}^\tau(2, \Delta) = A(h_1(\Delta, \tau) + h_2(\Delta, \tau)),$$

with

$$A = \frac{c_H^2 \nu^2}{(H + 1/2)^2},$$

$$h_1(\Delta, \tau) = \frac{1}{\tau^2} \int_{-\infty}^0 ((\Delta + \tau - s)^{H+1/2} - (\Delta - s)^{H+1/2} - (\tau - s)^{H+1/2} + (-s)^{H+1/2})^2 ds,$$

$$h_2(\Delta, \tau) = \frac{1}{\tau^2} \int_0^\Delta ((\Delta + \tau - s)^{H+1/2} - (\Delta - s)^{H+1/2})^2 ds.$$

We write $h_1(\Delta, \tau)$ under the form

$$\frac{1}{\tau^2} \Delta^{2H+2} \int_{-\infty}^0 \left(\left(1 + \frac{\tau}{\Delta} - s\right)^{H+1/2} - (1-s)^{H+1/2} - \left(\frac{\tau}{\Delta} - s\right)^{H+1/2} + (-s)^{H+1/2} \right)^2 ds.$$

Setting $\theta = \tau/\Delta$, we obtain

$$h_1(\Delta, \tau) = \Delta^{2H} f_1(\theta),$$

where

$$f_1(\theta) = \frac{1}{\theta^2} \int_{-\infty}^0 ((1+\theta-s)^{H+1/2} - (1-s)^{H+1/2} - (\theta-s)^{H+1/2} + (-s)^{H+1/2})^2 ds.$$

Similarly, we have

$$h_2(\Delta, \tau) = \Delta^{2H} f_2(\theta),$$

where

$$f_2(\theta) = \frac{1}{\theta^2} \int_0^1 ((1+\theta-s)^{H+1/2} - (1-s)^{H+1/2})^2 ds.$$

So

$$\hat{m}^\tau(2, \Delta) = A\Delta^{2H} (f_1(\theta) + f_2(\theta)).$$

Now remark that

$$\lim_{\theta \rightarrow 0} f_1(\theta) = (H+1/2)^2 \int_{-\infty}^0 ((1-s)^{H-1/2} - (-s)^{H-1/2})^2 ds$$

and

$$\lim_{\theta \rightarrow 0} f_2(\theta) = (H+1/2)^2 \int_0^1 (1-s)^{2H-1} ds.$$

Consequently,

$$\lim_{\theta \rightarrow 0} (f_1(\theta) + f_2(\theta)) = (H+1/2)^2 \frac{1}{c_H^2}.$$

Thus, when θ is small,

$$\hat{m}^\tau(2, \Delta) \sim \nu^2 \Delta^{2H}.$$

This means that the same scaling relationship as that associated to the spot volatility is approximately satisfied when considering implied volatilities with small enough times to maturity. Otherwise, one should add the multiplicative factor

$$f(\theta) = \frac{c_H^2}{(H+1/2)^2} (f_1(\theta) + f_2(\theta))$$

on the right hand side of the above relationship. This disrupts the scaling property and implies biased estimations for the Hurst parameter. We draw in Figure 9 the graph of the function f for $H = 0.04$.

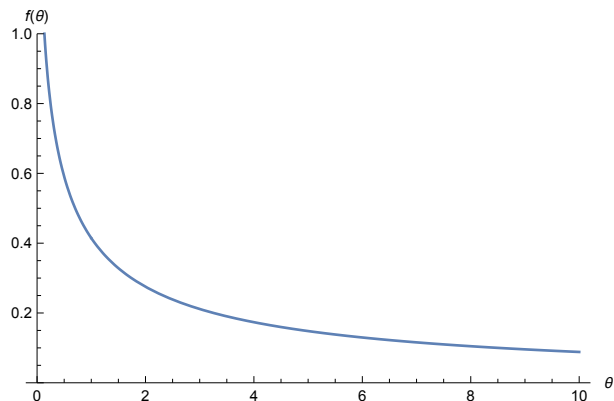


Figure 9: The function f for $H = 0.04$.

For fixed τ (as in Section 2), the function f is increasing with Δ . Therefore, when doing a regression analysis of the cloud of points with coordinates $(\log(\Delta), \log(\hat{m}^\tau(2, \Delta)))$, this implies an upward bias in the estimation of H due to a higher slope.

Acknowledgements

We thank Jim Gatheral and Olivier Scaillet for helpful comments and suggestions. We also thank the participants of the XVII Workshop on Quantitative Finance in Milano for fruitful discussions. Giulia Livieri acknowledges research support from the Scuola Normale Superiore.

References

- [1] C. Bayer, P. Friz, and J. Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6):887–904, 2016.
- [2] M. Bennedsen, A. Lunde, and M. S. Pakkanen. Hybrid scheme for Brownian semistationary processes. *arXiv preprint arXiv:1507.03004*, 2015.
- [3] M. Bennedsen, A. Lunde, and M. S. Pakkanen. Decoupling the short- and long-term behavior of stochastic volatility. *Available at SSRN 2846756*, 2016.
- [4] F. Black and M. Scholes. The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81:637–654, 1973.
- [5] Bloomberg. Introduction into the new Bloomberg implied volatility calculations. 2008.
- [6] E. Derman and I. Kani. Riding on a smile. *Risk*, 7(2):139–145, 1994.

- [7] B. Dupire. Pricing with a smile. *Risk*, 7(1):18–20, 1994.
- [8] O. El Euch, M. Fukasawa, and M. Rosenbaum. The microstructural foundations of leverage effect and rough volatility. *arXiv preprint arXiv:1609.05177*, 2016.
- [9] O. El Euch and M. Rosenbaum. The characteristic function of rough Heston models. *arXiv preprint arXiv:1609.02108*, 2016.
- [10] M. Forde and H. Zhang. Asymptotics for rough stochastic volatility and Lévy models. *Preprint*, 2015.
- [11] M. Fukasawa. Short-time at-the-money skew and rough fractional volatility. *Quantitative Finance, to appear*, 2016.
- [12] H. Funahashi and M. Kijima. Does the Hurst index matter for option prices under fractional volatility? *Annals of Finance*, pages 1–20, 2015.
- [13] J. Gatheral, T. Jaisson, and M. Rosenbaum. Volatility is rough. *Available at SSRN 2509457*, 2014.
- [14] H. Guennoun, A. Jacquier, and P. Roome. Asymptotic behaviour of the fractional Heston model. *Available at SSRN 2531468*, 2014.
- [15] P. S. Hagan, D. Kumar, A. S. Lesniewski, and D. E. Woodward. Managing smile risk. *Wilmott Magazine*, pages 84–108, 2002.
- [16] S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, 1993.
- [17] J. Hull and A. White. The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2):281–300, 1987.
- [18] T. Jaisson and M. Rosenbaum. Rough fractional diffusions as scaling limits of nearly unstable heavy-tailed Hawkes processes. *The Annals of Applied Probability*, 26(5):2860–2882, 2016.
- [19] A. Medvedev and O. Scaillet. Approximation and calibration of short-term implied volatilities under jump-diffusion stochastic volatility. *Review of Financial Studies*, 20(2):427–459, 2007.
- [20] J. Mühle-Karbe and M. Nutz. Small-time asymptotics of option prices and first absolute moments. *Journal of Applied Probability*, 48(4):1003–1020, 2011.
- [21] A. Neuenkirch and T. Shalaiko. The order barrier for strong approximation of rough volatility models. *arXiv preprint arXiv:1606.03854*, 2016.