

Estimating Rough Volatility Models

Bjørn Eraker *

Grigory Vilkov †

February 22, 2025

Abstract

Rough Volatility models are models where shocks to volatility are driven by Fractional Brownian Motion with Hurst index, H , less than one-half. We derive a MCMC based Bayesian estimator for the parameters in the RFSV model for which log volatility follows an Ornstein-Uhlenbeck driven by fractional Brownian motion. We estimate the Hurst exponent to be 0.3 with a tight posterior which rules out extreme roughness (H in the 0.04-0.15 range) reported in the extant rough volatility literature while also ruling out Brownian motion $H = 1/2$. The results differ from the extant rough volatility literature primarily because we employ a Bayesian filter for unobserved volatility which is less prone to sampling error than estimators based on Realized Volatility.

Preliminary and incomplete draft. Not for circulation.

*Wisconsin School of Business, Department of Finance. (bjorn.eraker@wisc.edu).

†Frankfurt School of Finance and Management. (vilkov@vilkov.net)

1 Introduction

In the recent decade, *rough volatility* models have received massive attention among practitioners and researchers in the quantitative finance literature. The literature, starting with Gatheral, Jaisson, and Rosenbaum (2018), argues that volatility sample paths exhibit erratic behavior at short frequencies consistent with *anti-persistent* short term negative autocorrelation between high-frequency volatility innovations. Much of the empirical evidence presented in this literature is based on the model’s ability to match features of realized volatility data, but the rough models have also been shown to fit SPX and VIX options implied volatility surfaces.

The rough volatility literature is at odds with the more traditional financial econometrics literature, in which one primary objective for the last three decades has been to develop models with high persistence (see for example classic work by Engle (1982) and Bollerslev (1986)) and even higher persistence (e.g., Baillie and Bollerslev (1994)) in the form of long-run memory. Indeed, Gatheral, Jaisson, and Rosenbaum (2018) propose a model they label rough fractional stochastic volatility (RFSV) by essentially just setting the Hurst parameter, H , in the long memory stochastic volatility model of Comte and Renault (1998) to a value less than one half. This introduces roughness. Gatheral, Jaisson, and Rosenbaum (2018) estimate the Hurst exponents to be something in the 0.1 to 0.2 range, providing strong evidence for the rough volatility hypothesis.

In this paper we examine the empirical validity of these two seemingly contradictory strands of the volatility literature. In doing so, we develop a full likelihood based estimation approach for analyzing historical asset returns data. Our approach uses Bayesian numerical (MCMC) procedure to sample from the posterior distribution of the RFSV model. A primary advantage of our approach is that we model the (potentially rough) volatility path as an unobserved process, as in Jacquier, Polson, and Rossi (1994), Harvey, Ruiz, and Shephard (1993), Eraker, Johannes, and Polson (2003), and others in the context of Markov models. This overcomes the problem of potentially falsely identified roughness due to measurement noise in volatility proxies such as realized variance (see Bolko, Christensen, Pakkanen, and Veliyev (2023)). By jointly estimating the speed of mean-reversion and the Hurts exponent in the RFSV model, we avoid potentially bi-




Figure 1: Top: log VIX, middle: gaussian OU process without roughness ($H = 1/2$). Bottom: RFSV model with $H = 0.3$. For the simulated processes, $\kappa = 0.0025$ and mean and variance are scaled to match sample mean and variance of log VIX.

ased estimates of either, as they are to some extent complimentary in shaping the autocorrelation function of volatility.

To further motivate the possibility of rough volatility, consider two simple pieces of evidence. First, Figure 1 plots the log VIX and two simulated sample paths for log volatility. The middle graph shows a standard log-normal OU process for volatility (no roughness), while the bottom graph shows log volatility simulated using the RFSV model. The visual impression left by the lower two graphs is that indeed, the rough path (bottom) *appears* to replicate the erratic behavior of log VIX more closely than the non-rough one (middle). Second, Figure 2 shows autocorrelations in first differences in daily log-VIX vs first differences in log volatility simulated




Figure 2: Top: autocorrelation of first difference in log VIX,. Bottom: autocorrelation of first difference in log volatility simulated from the RFSV model with $H = 0.302$ and $\kappa = 0.0084$.

from the RFSV model with $H = 0.3$. The picture shows that the RFSV model generates short-term statistically significant negative autocorrelation, as is also seen with log VIX. Notice that the first-order autocorrelation is actually greater in the model than in the data. The data also shows evidence of significant autocorrelation beyond lag 1, especially at lags 2 and 3.

1.1 Evidence of rough volatility

The rough volatility literature lists a number of other data-conforming features of rough models. First off, standard estimators for the Hurts exponent, H , tend to lie well below $1/2$ (Brownian motion). Gatheral, Jaisson, and Rosenbaum (2018) estimate it by matching moments of the

fBm to data over various sampling intervals and find estimates of about 0.14 for the SPX and 0.13 for Nasdaq indices. Gatheral et al. also match the term structure of skewness in SPX options data (defined as the derivative of the IV surface with respect to strike, K , for ATM options, $\text{skew}(\tau) := \partial IV(K, \tau) / \partial K$). Fukasawa (2011) shows that the skew has functional form $\tau^{H-1/2}$ in fractional models for small τ . Gatheral, Jaisson, and Rosenbaum (2018) use this result to calibrate the RFSV model can be made to match almost perfectly what appears to be a power-law decay in $\text{skew}(\tau)$ using $H = 0.1$. Euch, Gatheral, and Rosenbaum (2019) derive a rough version of the classic Heston (1993) model which notably has a semi-closed form solution for the characteristic function, allowing for standard Fourier inversion methods for computing option prices. Empirically, they show that their Rough Heston model delivers implied SPX volatilities that can be made to match data for two single days (Aug. 14, 2013 and May 19, 2017) with Hurst exponents $H = 0.12$ and 0.047 , respectively. Rømer (2022) provides a more comprehensive empirical analysis by calibrating various rough models to daily SPX options and find that Hurst exponents exhibit great amount of time-variation while averaging about 0.1. Brandi and Matteo (2022) apply the Generalized Hurst Estimator of Di Matteo, Aste, and Dacorogna (2003) to absolute returns and again report estimates of about 0.1 for various indices. Fukasawa, Takabatake, and Westphal (2019) derive an estimator akin to the Whittle estimator for fBm and estimate H to be 0.043 for S&P500 and similarly low values for other indices.

The literature on time-series evidence on full-scale, joint estimation of all the parameters in Rough Volatility models is limited. A notable exception is Bolko, Christensen, Pakkanen, and Veliyev (2023) who develop a GMM estimator for parameters in the RFSV model. They find estimates of the Hurts exponent that average about 0.024 when estimated from multiple equity indices and 0.043 for the S&P 500 (SPX), implying an extreme amount of roughness. Damian and Frey (2024) use an approximate Markov model to carry out particle filtering for fBm from returns data. Chong and Todorov (2024) derive a robust, non-parametric test of roughness and conclude that SPX volatility is rough but do not provide estimates of H .

Evidence of rough volatility is not entirely uncontroversial. Cont and Das (2023) provide evidence suggesting that standard methods for estimating the Hurts exponent can be biased downward so as to produce false evidence of roughness due to estimation noise in realized volatility.

Noteworthy, they find that even with $H = 1/2$ (no roughness), H can be estimated to be 0.13. They conclude that “... results suggest that the origin of the roughness observed in realized volatility time-series lies in the estimation error rather than the volatility process itself” (see abstract). Their evidence is not an indictment of rough volatility models per se, but rather an indictment of evidence based on realized volatility, which is artificially noisy due to estimation error, and thus motivates further development of methods, such as the one proposed here, that utilize filtering rather than realized volatility estimates, to conduct econometric analysis. Bolko, Christensen, Pakkanen, and Veliyev (2023) also note that the use of RV induces *illusiv e roughness* as estimation noise contaminates estimated sample paths, leading to biased estimates of the Hurst exponent.

1.2 Likelihood inference

Likelihood inference for processes driven by fractional Brownian motion is complicated because the processes are not Markovian. This means that one cannot factor the likelihood function as a product of conditional distributions, as for example with an AR(1) process for log volatility, Y_t , which likelihood function can be written $\prod_{t=1}^T p(Y_t \mid Y_{t-1})$ which leads to a conditional posterior of the form $p(Y_t \mid Y_{t-1}, Y_{t+1}, \text{Data}) \propto p(Y_t \mid Y_{t-1})p(Y_{t+1} \mid Y_t)p(\text{Data} \mid Y_t)$. Instead, for processes driven by fBM, the likelihood function will depend on the entire history $Y_t, t = 0, \dots, T$ which in turn requires on the order of T^2 operations to compute. It also invalidates one-step-at-a-time MCMC updating schemes that were the basis for the posterior sampling schemes in Jacquier, Polson, and Rossi (1994), Harvey, Ruiz, and Shephard (1993) and Eraker, Johannes, and Polson (2003), among others. On the other hand, one-at-a-time updating typically leads to slowly converging MCMC chains, and thus typically requires longer sample sizes to accurately estimate the posterior distributions. In this paper, we demonstrate that it is possible to draw *the entire path* of unobserved latent log-volatilities in a single MCMC step. This is made possible by utilizing the known (auto) covariance structure for any $t \in (0, T)$ for fBM which can be used to find the entire auto covariance matrix of the (log) volatility process. The approach is still $O(T^2)$ but made computationally tractable because it requires only one single Metropolis-

Hastings MCMC step. Thus, this approach is potentially useful, even if estimating Markovian models as in Jacquier, Polson, and Rossi (1994), Harvey, Ruiz, and Shephard (1993), and others.

1.3 Summary of Findings

We estimate the Hurst exponent to be almost exactly 0.3. The marginal posterior distribution of H is very concentrated and puts almost no mass on values below 0.29, suggesting that volatility, while rough, is not nearly as rough as some estimates based on Realized Variance and derivatives data suggest. The posterior is also well bounded away from $1/2$, suggesting that within the parametric framework of the RFSV model, Brownian motion can be ruled out as a driver of volatility innovations. The estimated model produces an autocorrelation pattern in log-volatility which decays slower than an OU process driven by Brownian motion, indicating that $H = 0.3$ also generates long-run dependence. Simulations show that off-the-shelf estimators of the Hurst exponent produce severely downwardly biased estimates when applied to realized variance data, but can be very accurate when applied to the actual true spot volatility but since spot volatility is unobserved in practice, this is an infeasible estimator. The implication of this is that filtering approaches, such as the one used here, provide superior inference relative to off-the-shelf estimates based on noisy estimates of volatility, and suggest an econometric advantage for methods that jointly estimate the model parameters and spot-volatility, and in particular, implement Bayesian filtering for the unobserved volatility process.

The remainder of the paper is organized as follows: Section 2 outlines the econometric method. Section 3 presents the estimation results and discusses the model implications for modeling returns, volatility forecasting, and estimation from realized variance. Section 4 summarizes the findings.

2 Rough Volatility Models

2.1 Fractional Brownian Motion

Fractional Brownian motion is defined as a process,

$$W_t^H = \int_0^t K(s, t) dW_s \quad (1)$$

where $K(s, t)$ is a kernel function and W_s is a standard Brownian motion. The Riemann-Luiville representation has the kernel

$$K(s, t) = \frac{1}{\Gamma(H + 1/2)} \int_0^t (t - s)^{H-1/2}, \quad (2)$$

although alternative representations are constructed in the literature. Note that eqn. (1) can be thought of as a continuous time analogue of a discrete-time moving average process $Y_t = \sum_{s=0}^t \varphi_s \epsilon_s$ with MA(s) parameters φ with power-law decay.

The autocovariance function of fBm is given by

$$G(H) = \mathbb{E}[W_t^H W_s^H] = \frac{1}{2} (|t|^{2H} + |s|^{2H} - |t - s|^{2H}) \quad (3)$$

which reduces to

$$G(1/2) = \min(s, t) \quad (4)$$

in the case of Brownian motion. Notice that this implies that the likelihood function for discretely observed realizations of fBm can be constructed from the fact that the joint distribution of $W = \{W_t\}_{t=0, \dots, T}$ is

$$W \sim N(0, G(H)) \quad (5)$$

which we will utilize to compute the likelihood function for the RFSV model.

In the case $H > 1/2$ the fBm exhibits long-memory and if $H < 1/2$ it exhibits *rough paths*. Roughness is exemplified in Figure 1 and generates negative autocorrelation in the increments

$W_t^H - W_{t-1}^H$ which translates into short-term negative autocorrelation for the increments of the log volatility process in the RFSV model, as seen in Fig. 2.

Define the matrix C to be the lower triangular Cholesky decomposition of $G(H)$ such that $G = CC'$. We can then simulate discrete realizations of W^H through $W = C\eta$ where η is a T length vector of $N(0, 1)$ s.

2.1.1 Path-dependence

Notice that while the originating Brownian motion W obviously has independent increments, W^H does not. When $H \neq 1/2$ the an increment $W_{t+\Delta} - W_t$ can be predicted from past realizations W_s , $s < t$. It is also not Markovian, as the increments over some period of time depend on the entire past of the process. This complicates simulation as one of the standard methods for simulating fBm is by Cholesky decompose $G(H)$ - an $O(T^3)$ operation. This paper uses the Cholesky decomposition as a primary means of constructing an MCMC sampler, implying high computational demands.

2.2 The RFSV models

As a baseline model we consider the following

Model 1

$$d \ln P_t = \mu dt + \sigma_t \rho dW_t + \sigma_t \sqrt{1 - \rho^2} dB_t \quad (6)$$

$$\sigma_t^2 = e^{a+bY_t} \quad (7)$$

$$dY_t = -\kappa Y_t dt + dW_t^H \quad (8)$$

dubbed RFSV (Rough Fractional Stochastic Volatility) by Gatheral, Jaisson, and Rosenbaum (2018). We allow the originating Brownian motions W and B to be correlated,

$$\text{Corr}(dB_t, dW_t) = \rho \quad (9)$$

so as to generate a “leverage effect” which is crucial in capturing the negative correlation in returns and spot-volatility. Thus, the parameter vector is $\Theta = \{a, b, \rho, \kappa, H\}$. Note that the model is observationally equivalent to a model with $a = 0$ and $b = 1$ with spot volatility

$$dY_t = \kappa(\theta - Y_t)dt + \eta dW_t^H, \quad (10)$$

where θ represents the average volatility and η the volatility-of-volatility. The model in eqns. (6) - (8) is parameterized so as to improve efficiency of the numerical Bayesian computations.

Model 2. Quasi equilibrium

The second formulation is

$$d \ln P_t = \mu dt + \beta d\sigma_t^2 + \sigma_t dB_t, \quad (11)$$

$$\sigma_t^2 = e^{a+bY_t}, \quad (12)$$

$$dY_t = -\kappa Y_t dt + dW_t^H. \quad (13)$$

Here, the term βdY_t replaces the leverage term ρ in Model 1 is zero and the log-volatility shocks dY_t directly enter the price dynamics in (11). This is motivated by equilibrium price dynamics in models such as those by Bansal and Yaron (2004), Eraker and Wu (2017) and others where an equilibrium price is obtained as $\ln P_t = \ln D_t + \beta \sigma_t^2 + \dots$ where D_t is a “fundamental” such as dividends, earnings, or some process that generates a terminal wealth payoff. Notice that Model 2 introduces a small fractional component into the log-price of the asset itself as here the price process depends directly on the log-volatility process, Y . Thus, the price process is not a semi-martingale.

The main objective of our paper is to compute the posterior distribution of Θ from a discrete sample of return, $\ln R_t = \ln P_t - \ln P_{t-1}$. By jointly estimating the parameters of the model we mitigate potential biases that result from estimating the Hurts exponent using “off the shelf” methods such as applied to noisy approximations to spot volatility, such as realized volatility.

3 Econometric Method

3.1 Simulating the RFSV Model

First note that the unconditional distribution of the latent log-volatility, Y , in the RFSV model in eqns. (6) - (8) can be written

$$Y = CdW \quad (14)$$

where dW is a T vector of discrete time increments to Brownian motion. The matrix $C = C(\kappa, H)$ is the matrix root of

$$V = A(\kappa)G(H)A(\kappa)' \quad (15)$$

where A is the matrix

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ b-1 & 1 & 0 & 0 & \cdots & 0 \\ b^2-b & b-1 & 1 & 0 & \cdots & 0 \\ b^3-b^2 & b^2-b & b-1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & 0 \\ b^{T-1}-b^{T-2} & b^{T-2}-b^{T-3} & \cdots & b^2-b & b-1 & 1 \end{bmatrix} \quad (16)$$

with $b = e^{-\kappa}$ and $G(H)$ is the autocovariance matrix of fBm given in eqn. (3).

If B is a lower triangular matrix such that $BB' = G$ then

$$C(\kappa, H) = A(\kappa)B(H) \quad (17)$$

and simulation of the model can be carried out using (14).

3.2 The Likelihood Function

Let $\Theta = \{\mu, \rho, a, b, \kappa, H\}$ and let R and Y denote the entire sample of returns and log-volatilities, respectively.

The likelihood function is

$$p(\Theta \mid R, Y) \propto p(R \mid \rho, a, b, Y)p(Y \mid \kappa, H). \quad (18)$$

To compute this $p(R \mid \rho, a, b, Y)$ for a given Y , compute

$$dW = C^{-1}Y \quad (19)$$

such that

$$z_t = \ln R_t - \mu - \rho\sigma_t dW_t \quad (20)$$

has distribution $N(0, x_t^2)$ where

$$x_t^2 = \sigma_t^2(1 - \rho^2) = e^{a+bY_t}(1 - \rho^2) \quad (21)$$

This gives the conditional log-likelihood

$$\begin{aligned} \ln p(R \mid Y, \mu, \rho, a, b) = & -\frac{1}{2}T(1 - \rho^2) - \frac{1}{2} \sum_t (a + bY_t) \\ & - \frac{1}{2} \sum_t (\ln R_t - \mu - \rho\sigma_t dW_t)^2 \frac{e^{-a-bY_t}}{1 - \rho^2} \end{aligned} \quad (22)$$

for the returns conditional upon Y while the likelihood for the log-variance Y process is

$$\begin{aligned} \ln p(Y \mid \kappa, H) &= -\frac{1}{2} \ln |V| - \frac{1}{2} Y' V^{-1} Y \\ &= -\frac{1}{2} \ln |V| - \frac{1}{2} dW' dW \end{aligned} \quad (23)$$

The conditional posterior for the parameters is given by

$$p(\Theta \mid R, Y) \propto p(R \mid Y, \mu, \rho, a, b)p(Y \mid \kappa, H) \quad (24)$$

by Bayes theorem.

To sample from the posterior, we draw Θ and Y in two alternate Metropolis within Gibbs sampling steps.

3.3 Log-volatility step

To sample Y note that we have

$$p(Y \mid R, \Theta) \propto p(R \mid Y, \mu, \rho, a, b)p(Y \mid \kappa, H) \quad (25)$$

which requires the computation of $V^{(g-1)} = V(\kappa^{(g-1)}, H^{(g-1)})$ - an $O(T^2)$ operation. Let c be a small constant. We draw a proposal

$$Y_p = Y^{(g-1)} + x + cC^{(g-1)}\eta \quad (26)$$

where $C^{(g-1)}$ is the Cholesky decomposition $C^{(g-1)}C^{(g-1)'} = V^{(g-1)}$ of $V^{(g-1)}$ and η is a T length vector of $N(0, 1)$ s and x is a random scalar drawn from $N(0, \sigma_x^2)$ for some tuning parameter σ_x . The addition of the scalar x to the vector of proposals ensures that the mean of the proposed value Y_p can move independently of the draws for η which improves mixing in the MCMC chain.

The proposal in (26) is accepted with probability $\alpha = p(Y_p \mid R, \Theta)/p(Y \mid R, \Theta)$ and is referred to as a random walk sampler in the MCMC literature. The parameter c , which scales the variance of the proposals, is typically set to some low value in the 0.01 to 0.03 range. This implies that the volatility sampler will gradually perturbate the current estimate in small random increments around its current value. Remarkably, this sampler works in any dimension T , primarily because the covariance structure of the posterior is dominated by the covariance structure implied by V . For computational considerations, as V and thus C , which depend on κ and H , are computationally expensive, it is optimal to sample a certain number of paths of Y before updating the parameters. This ensures better mixing in the MCMC chain.

3.4 Parameter Step

We employ the Metropolis-Hastings algorithm to draw $\Theta^{(g)}$ in the g th iteration of a Metropolis within Gibbs scheme. This entails drawing a proposal, say Θ_p from a proposal density $q(\Theta)$ and accept/reject this with probability,

$$\alpha = \min \left[\frac{p(\Theta_p | R, Y^{(g)})q(\Theta^{(g-1)})}{p(\Theta^{(g-1)} | R, Y^{(g)})q(\Theta_p)}, 1 \right]. \quad (27)$$

Since the conditional posterior, $p(\Theta_p | R, Y^{(g)})$ is time-consuming to compute, we derive a sampling scheme designed to accept or reject the parameter vector Θ in a single step. To do this, we carefully tailor a joint proposal density derived from a sequence of conditional distributions that incorporate information about the mode and dispersion in the conditional posterior. In many cases, this can be accomplished by interpreting the parameters as linear regression or correlations between observables. For example, the conditional posterior for κ is approximately normal with mean and variance given by the standard expressions for the slope and its posterior standard deviation in a regression of $dY = \{Y_t - Y_{t-1}\}$ on $Y = \{Y_t\}$. Similarly, ρ is a correlation between the Brownian shocks which can be approximated by the standard expressions for the Pearson correlation coefficient computed from the standardized returns and the estimated Brownian shocks to the log volatility process, Y .

Another noteworthy feature of the model is that the parameters a, b and ρ depend on $p(R | Y, \mu, \rho, a, b)$ but not $p(Y | \kappa, H)$. But $p(R | Y, \mu, \rho, a, b)$ can be computed in $O(T)$ operations, which is computationally quick. For this reason, it is then possible to update these parameters more frequently than updating κ and H which is computationally costly. Essentially, we can create a “sub-chain” of MCMC draws for a, b and ρ that when run for, say N sub-iterations, creates draws that are near independent.




Figure 3: Posterior distributions for parameters in the RFSV model using SPY data, 1993-2025 (8060 obs).

4 Empirical Results

4.1 Parameter Estimates

Table 2 and Figure 3 present the estimation results. The parameters a and b determine the level and volatility of the volatility process, respectively. The posterior mean of a is -0.0537 corresponding to an average daily standard deviation of log returns of $\exp(-0.0537/2) = 0.974$ which is in line with the average SPX volatility of about 1% per day. The leverage parameter, ρ , is estimated to be about -0.7 which again is consistent with many prior studies that find a strong negative correlation between innovations in volatility and returns. The mean-reversion parameter, κ , is estimated to be about 0.01 . If H had been $1/2$ (Brownian motion), this estimate

Table 1: Parameter Estimates - RFSV Model

The table reports posterior means and standard deviations for the RFSV model based on daily SPY closing prices from 1993-2025 (8060 obs) and subperiods 1993-2000 (XX obs), 2000-2010 (YY obs),

1993-2025					
	a	b	ρ	κ	H
Mean	-0.261	0.376	-0.702	0.00835	0.304
Std.	0.0639	0.0158	0.0275	0.00127	0.0073
1993-2010					
Mean	-0.241	0.376	-0.677	0.00324	0.237
Std	0.156	0.0286	0.0329	0.00103	0.0152
2010-2020					
Mean	-0.062	0.299	-0.732	0.00326	0.285
Std	0.293	0.0523	0.0534	0.00183	0.0463
2020-2025					
Mean	-0.00464	0.377	-0.722	0.0216	0.356
Std	0.133	0.0542	0.0481	0.00719	0.0455

would have implied an autocorrelation for log-volatility of about $1 - 0.01 = 0.99$. Further analysis of the autocorrelation properties of the model will be presented below. Finally, the most interesting parameter from the point of view of this study is the Hurst exponent H which is estimated to be 0.3 almost exactly. The posterior distribution for H has very little mass above 0.3 and almost no mass at all below. This is different from when running the estimator on simulated data, for which the posterior tends to put mass on both sides of the true H used to generate artificial data. The posterior distribution for H seen in Figure 3 is highly concentrated, and right-skewed to put some mass on values as high as 0.32 but rules out many of the estimates reported in the rough volatility literature, as referenced in the introduction. It also rules out

Brownian motion ($H = 1/2$). Thus, the results presented here provide almost overwhelming evidence against both the estimates found in the rough volatility literature, and the time-series literature based on Brownian motion.

Table 2 also contains estimates based on subperiods. The results show that using the earlier sample periods leads to lower estimates of H than do the more recent 2020 to 2025 sample period. Note that the latter period uses only about 1300 data points for estimation which also significantly increases the posterior standard deviation.

4.2 Properties of the model




Figure 4: Histograms of SPY log-returns vs log-returns simulated from the estimated RFSV model.

Figure 4 shows the vs model frequency histograms of returns. There actual data clusters more frequently around zero and slightly above than do the data simulated from the model. The model captures well the tails of the unconditional distribution.






Figure 5: Autocorrelations of the estimated RFSV model using $H = 0.3$ and $H = 0.5$.

Figure 5 shows the autocorrelations function (ACF) of log volatility in the RFSV model using $H = 0.3$ and $H = 0.5$. Rough volatility ($H = 0.3$) generates significantly slower decay in the ACF than does Brownian motion ($H = 0.5$).


To investigate the dynamic properties of the model further, it is useful to remember that fractional Brownian motion is path-dependent. For this reason, it is not useful to think about how a single shock affects the behavior of the process as is the case in Markov models. This implies that we cannot simply study an impulse response function to learn how a single shock will impact the process going forward, but rather that we need to condition on the entire past history of the process. To this end, we compute expected values, or forecasts, of Y conditional upon realized sample paths that are deliberately chosen to have either 1) low or high starting values, or 2) display positive or negative average increments such as to suggest the appearance of a “trend” despite being random draws that just happen to have, on average, positive or negative increments.



(a) Forecast $Y_t > 0$ following a negative shock



(b) Forecast $Y_t < 0$ following a positive shock



(c) Long-term forecasts.

Figure 6: RFSV forecasts for log-volatility. The blue areas represent the observed paths of Y , pink represents forecasted values.

Figure 6 illustrate the effect of path-dependence and roughness on forecasts of log-volatility, Y . In the rightmost figure (a) the Y process is above its mean (0) for most of the time prior to the forecast made at day 250 and shown in the pink area of the graph. It illustrates that that a negative shock on the day prior to the forecast implies short-term reversal to the upside, illustrated by the hump-shaped forecast over the next near-term. Likewise, graph (b) illustrates the same effect when the process is below its mean and experiences a positive shock the day prior to making the forecast, as now the near-term forecast is for a lower (log) volatility while the longer-term is dominated by reversal to the mean. Graph (c) shows the effect of a positive (blue) and negative (black) trajectory of Y prior to the forecast date. In both cases, the expected

values for about 400 days out overshoot the mean of zero. This behavior illustrates the path-dependent behavior of fBm as the positive (negative) past increments propagate into a forecast that exceeds (less than) the unconditional mean of the process.

4.2.1 Realized Volatility




Figure 7: Simulated spot RFSV volatility vs Realized Variance.

To gauge the impact of estimation noise in off-the-shelf estimators of the Hurst index, we simulate high-frequency data from the RFSV model using estimated parameters in Table 2. The data are simulated using $\Delta = 1/66$ corresponding to typical 5-minute return sampling intervals used in the literature. Figure 7 depicts the a two-year simulated sample path of the RFSV model with $H = 0.3$ and RV. The plot illustrates clearly the problem with RV: while centered reasonably around the true spot variance, it is extremely noisy.

To gauge the possible impact on the estimation of H from RV, Table 2 reports results from applying two “off-the-shelf” estimators for the Hurst exponent to simulated log-volatility data using the RFSV model with parameters as reported in Table 2 but with $H = 0.3$ exactly. The

Table 2: Estimates of H from log-volatility and RV

The table reports estimates of the Hurst exponent H estimated using the Whittle estimator and the Generalized Hurst Estimator (GHE) of T. Di Matteo et al (2003). GHE(q) denotes which moment is used in the GHE estimator. The estimators are applied to log-volatilities, Y , simulated from the RFSV model at the 5-minute interval and daily realized variance (RV) computed from 5-minute returns over 500 days. Results based on 10,000 Monte-Carlo experiments.

Data Estimator	Y Whittle	Y GHE(2)	RV Whittle	RV GHE(1.5)	RV GHE(2)	RV GHE(3)
Mean	0.252	0.307	0.010	0.029	0.028	0.028
Std	0.024	0.004	0.000	0.018	0.018	0.019
$\Pr(\hat{H} > 0.3)$	0.028	0.956	0.000	0.000	0.000	0.000

two rightmost columns show that if one were to observe the actual spot volatility process Y , these estimators do a good and very good job, respectively, of identifying it from data. The Generalized Hurst Estimator of Di Matteo, Aste, and Dacorogna (2003) almost exactly identifies the true value of H with an average estimate of 0.307 although 95% of the simulated values are above the truth. The Whittle estimator is downward biased by about 0.05, giving an average estimate of 0.252. Using RV gives severe downward bias in the estimate of H for both estimators, typically producing values close to zero and in some cases hitting the lower bound for the estimator (0.01). This is consistent with the observed “illusory roughness” depicted for RV in Figure 7 and also with the evidence in Cont and Das (2023) and Bolko, Christensen, Pakkanen, and Veliyev (2023). Overall, the evidence suggests that RV is unstable and would need to be smoothed or filtered in some way to be reliably used to recover evidence of rough volatility.

4.3 Re-examining VIX

There is some theoretical justification for using VIX or $\log(\text{VIX})$ as a proxy for spot volatility: in affine single models, such as Heston (1993), squared VIX is a linear function of spot volatility. As such the dynamic properties of spot volatility, including speed of mean reversion, and the Hurst index can be estimated under the joint null of $H = 1/2$ and the model being correct.




Figure 8: Hurst exponents estimated from yearly log VIX. The estimates are based on the GHE method of Di Matteo et al (2003).

In rough models, the path-dependency will imply that the one-to-one map between spot and implied volatility breaks down. With this caveat in mind, Figure 8 plots year-by-year estimates of H using log VIX data. The Figure shows significantly less convincing evidence of roughness than what is reported in the rough volatility literature with point estimates averaging about 0.4. When using the the whole 1993-2025 sample H is estimated to be about 0.4 with a standard error of 0.023. This rules out $H = 1/2$ in the population but it also rules out 0.3 estimated using our Bayesian estimator. This evidence is also consistent with that shown in Fig. 2 in the introduction where we see that differenced log VIX has less negative first-order autocorrelation than does data simulated from the RFSV model.

5 Concluding Remarks

This paper examines rough volatility and develops a full likelihood-based approach for estimating the model parameters in the RFSV model. This allows us to examine the claim in the *rough volatility* literature that spot volatility exhibits roughness in that the Hurst exponent in the fractional Brownian motion assumed to drive innovations in spot volatility is below $1/2$ (Brownian motion). The evidence suggests that indeed, volatility is rough, but not nearly as rough as suggested by many authors in the rough literature. Using our largest sample of daily SPX returns, we find that the Hurst exponent is estimated very precisely to 0.3. For subperiods, the estimates vary from 0.23 to 0.36 suggesting that perhaps recent data is less suggestive even of roughness than data from the earlier sample periods. We also rule out Brownian motion ($H = 1/2$) as a driving source of uncertainty. Thus our evidence is consistent with roughness, just much less than what is typically argued in the rough volatility literature.

A primary source of the difference in the results reported here relative to the extant rough literature is the use of a Bayesian filter for spot volatility. The filter will produce smoothed estimates of the latent volatility, alleviating the problems reported in Cont and Das (2023) and Bolko, Christensen, Pakkanen, and Veliyev (2023) in applying off-the-shelf estimators of the Hurst exponent on RV data. The intuition behind this can be illustrated in Fig. 7. The RV estimator relies entirely on non-overlapping subsampling of noisy returns to form an estimate RV_t say, but ignores the information about the spot-volatility prior to date t , thus ignoring that the continuity of the volatility process is unlikely to lead to a substantial change in the process over a small sampling interval. Bayesian filtering incorporates this information through the probability model for the volatility process. In Markov models it suffices to condition on Y_{t-1} to filter Y_t but in non-markovian models such as the one considered here, Bayesian MCMC filtering implies that the conditional distribution of the process at a single time point is estimated conditional upon the entire sequence $Y_s, s \in (0, T) \setminus t$ and also that distant observations are informative about the process at time t .

Appendix

A Computational Considerations

The likelihood function requires the inverse of the $T \times T$ matrix $V = AGA'$. Let $BB' = G$ be the Cholesky decomposition of G . We then have

$$V^{-1} = (A')^{-1}(B^{-1})'B^{-1}A^{-1} := C^{-1}(C^{-1})'.$$

Since $A^{-1} = A^{-1}(\kappa)$, $B^{-1} = B^{-1}(H)$ and $C^{-1} = C^{-1}(\kappa, H)$ we can reduce the computational burden of re-computing these matrices by computing them over a discrete grid of values for κ and H allowing many of the computations required for likelihood inference outlined here to avoid having to recompute matrix inverses and Cholesky decompositions (an $O(T^3)$ operation) of large T by T matrices in for each iteration of the MCMC sampler, effectively reducing the computational burden to involve only matrix multiplications.

Generally, C^{-1} can be stored as a band-diagonal matrix which eases the computation $dW = C^{-1}Y$ but C cannot as its off-diagonal elements typically do not converge to zero sufficiently fast for a band-diagonal approximation to be accurate.

The empirical results reported here were computed using Matlab, which is computationally fast for matrix operations, running on various computers including an Intel core based Linux server with NVIDIA T2 GPUs chips and 500GB of internal memory. This allows the use of 'gpuArray' functionality in Matlab which parses large matrix multiplications to the GPUs. With $T = 8060$ observations a single iteration of the MCMC sampler, consisting of 100 draws of the latent log-volatility Y and 50 draws of a, b and ρ for each update of κ and H , takes between one and two seconds leading to a computing time between one and two days to sample 100,000 posterior draws, which is typically enough to obtain reasonably accurate estimates of the posteriors. Smaller sample sizes are considerably faster.

Figure 9 depicts the MCMC output for the parameters (trace plots) for three independent runs of the sampler each started near the mode of the likelihood function as established by a




Figure 9: Trace plots of the MCMC sampler.

previous MCMC sampler. They reveal relatively strong serial dependence in the parameters b , ρ and H , adding computational burden to the estimation.

References

- Baillie, R. T., and T. Bollerslev, 1994, Long memory in the forward premium, *Journal of International Money and Finance* 13, 565–571.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance* 59, 1481–1509.
- Bolko, Anine E., Kim Christensen, Mikko S. Pakkanen, and Bezirgen Veliyev, 2023, A GMM approach to estimate the roughness of stochastic volatility, *Journal of Econometrics* 235, 745–778.
- Bollerslev, T., 1986, Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics* 31, 307–327.
- Brandi, Giuseppe, and T. Di Matteo, 2022, Multiscaling and rough volatility: an empirical investigation, *arXiv preprint arXiv:2201.10466*.
- Chong, Carsten H., and Viktor Todorov, 2024, A nonparametric test for rough volatility, *Working Paper, Northwestern*.
- Comte, F., and E. Renault, 1998, Long Memory in Continuous-Time Stochastic Volatility Models, *Mathematical Finance* 8, 291–323.
- Cont, Rama, and Purba Das, 2023, Rough volatility: fact or artifact?, *Working Paper, Oxford*.
- Damian, Camilla, and Rüdiger Frey, 2024, Detecting rough volatility: a filtering approach, *Quantitative Finance* 24, 1493–1508.
- Di Matteo, T., T. Aste, and M.M. Dacorogna, 2003, Scaling behaviors in differently developed markets, *Physica A: Statistical Mechanics and its Applications* 324, 183–188 Proceedings of the International Econophysics Conference.
- Engle, Robert F., 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987–1007.

- Eraker, Bjørn, Michael. J. Johannes, and Nicholas G. Polson, 2003, The Impact of Jumps in Returns and Volatility, *Journal of Finance* 53, 1269–1300.
- Eraker, Bjørn, and Yue Wu, 2017, Explaining the Negative Returns to VIX Futures and ETNs: An Equilibrium Approach, *Journal of Financial Economics* 125, 72–98.
- Euch, Omar El, Jim Gatheral, and Mathieu Rosenbaum, 2019, Roughening Heston, *Risk* May 2019, 84–89.
- Fukasawa, M., 2011, Asymptotic analysis for stochastic volatility: Martingale expansion, *Finance and Stochastics* 15(4), 635–654.
- Fukasawa, Masaaki, Tetsuya Takabatake, and Rebecca Westphal, 2019, Is Volatility Rough?, *arXiv preprint arXiv:1905.04852*.
- Gatheral, Jim, Thibault Jaisson, and Mathieu Rosenbaum, 2018, Volatility is rough, *Quantitative Finance* 18, 933–949.
- Harvey, A., E. Ruiz, and N. Shephard, 1993, Multivariate Stochastic Variance Models, *Review of Economic Studies* 61, 247–264.
- Heston, Steve, 1993, Closed-Form Solution of Options with Stochastic Volatility with Application to Bond and Currency Options, *Review of Financial Studies* 6, 327–343.
- Jacquier, E., N. G. Polson, and P. Rossi, 1994, Bayesian analysis of stochastic volatility models, *Journal of Business and Economic Statistics* 12, 371–389.
- Rømer, Sigurd Emil, 2022, Empirical analysis of rough and classical stochastic volatility models to the SPX and VIX markets, *Quantitative Finance* 22, 1805–1838.