Journal of Mathematics and Modeling in Finance (JMMF)
Vol. 6, No. 1, Winter & Spring 2026
Research paper

Volatility Modeling via EWMA-Driven Time-Dependent Hurst Parameters

Athipatla, Jayanth[1]

[1] Department of Mathematics, University of Nebraska at Omaha, Nebraska, United States
jathipatla@unomaha.edu

Abstract:
We introduce a novel rough Bergomi (rBergomi) model featuring a variance-driven exponentially weighted moving average (EWMA) time-dependent Hurst parameter H_t, fundamentally distinct from recent machine learning and wavelet-based approaches in the literature. Our framework pioneers a unified rough differential equation (RDE) formulation grounded in rough path theory, where the Hurst parameter dynamically adapts to evolving volatility regimes through a continuous EWMA mechanism tied to instantaneous variance. Unlike discrete model-switching or computationally intensive forecasting methods, our approach provides mathematical tractability while capturing volatility clustering and roughness bursts. We rigorously establish the existence and uniqueness of solutions via rough path theory and derive martingale properties. Empirical validation in diverse asset classes including equities, cryptocurrencies, and commodities demonstrates superior performance in capturing dynamics and out-of-sample pricing accuracy. Our results show significant improvements over traditional constant-Hurst models.

Keywords: Stochastic, Volatility, Hurst, Rough.
Classification: 60G22, 91G20

1 Introduction

The modeling of volatility dynamics in financial markets has evolved substantially over the past several decades, driven by the need to capture empirical stylized facts such as volatility clustering, long memory, and non-Markovian dependence. Classical stochastic volatility models, including Heston-type diffusions, improve upon constant-volatility assumptions, but remain fundamentally Markovian and struggle to reproduce the observed roughness of volatility paths at short time scales.

The rough volatility paradigm, initiated by Gatheral et al. 2014 [4] and formalized through the rough Bergomi model (rBergomi) by Bayer et al. 2016 [1] represents a major advance in this direction. By modeling logarithmic volatility as a fractional process with the Hurst parameter $H < 1/2$, rough volatility models successfully capture the steep implied volatility skews and persistent memory observed in high-frequency financial data. Empirical studies consistently estimate the Hurst

[1] Corresponding author
Received: 23/09/2025 Accepted: 27/12/2025
https://doi.org/10.22054/jmmf.2026.88385.1218

parameter to lie in a narrow range around 0.1, suggesting an intrinsic roughness of volatility across asset classes.

Despite their empirical success, standard rough volatility models assume a constant Hurst parameter throughout time. This assumption imposes a homogeneous roughness structure that is in conflict with observed market behavior, particularly during periods of stress or regime transitions. Empirical evidence indicates that volatility roughness itself varies between market conditions, with bursts of increased irregularity during crises and calmer dynamics during stable periods. A constant Hurst parameter therefore limits the ability of rough volatility models to adapt to evolving market regimes.

Recent work has begun to explore time-varying roughness to address this limitation. Shah [5] proposes a machine learning-based framework that forecasts Hurst parameters and switches discretely between rBergomi and classical stochastic volatility regimes for American option pricing. Orefice [6] and Webb [14] employ multifractal and wavelet-based techniques to infer changing rougness from high-frequency data, linking the inferred Hurst dynamics to implied volatility features such as the at-the-money skew. While these approaches provide valuable empirical insights, they often rely on computationally intensive estimation procedures, discrete regime switching, or statistical constructions that are difficult to embed into a continuous-time pricing framework with rigorous theoretical guaranties.

The present paper introduces a fundamentally different approach to time-varying roughness. We propose a variance-driven, exponentially weighted moving average (EWMA) time-dependent Hurst parameter embedded directly within the rough Bergomi framework. Rather than forecasting the Hurst parameter externally or switching between discrete regimes, we allow roughness to evolve endogenously as a continuous function of past volatility. Recent variance observations exert greater influence through the EWMA mechanism, while older information decays exponentially, yielding a smooth and interpretable roughness path that adapts naturally to changing market conditions.

From a theoretical perspective, our model is formulated within a unified rough Volterra framework grounded in rough path theory. The time-dependent Hurst parameter enters the kernel in an adapted, non-anticipative manner, preserving mathematical tractability. We establish the existence and uniqueness of solutions, derive Gaussian properties of the variance driver, and analyze martingale conditions under both correlated and uncorrelated settings. Unlike many empirical models of time-varying roughness, our construction admits a rigorous continuous-time formulation suitable for derivative pricing and risk-neutral valuation.

From a practical point of view, the proposed rough Bergomi model driven by EWMA offers significant computational advantages. It avoids the overhead of machine learning–based forecasting and the discontinuities associated with regime-switching models, while remaining straightforward to implement via standard Euler–Maruyama discretization. Empirical results in equities, cryptocurrencies and

commodities demonstrate consistent improvements over constant-Hurst rBergomi and Heston models in distributional fit, volatility autocorrelation structure, and out-of-sample option pricing accuracy.

This paper contributes to the literature by introducing a fundamentally different approach, a time-dependent Hurst parameter based on EWMA-based on variance within a unified rBergomi RDE framework. Our methodology distinguishes itself from existing approaches in several key aspects: computational simplicity compared to ML-intensive forecasting methods, continuous evolution without discrete regime switching, direct variance dependence rather than indirect wavelet or multifractal mechanisms, and rigorous mathematical foundation via rough path theory.

Our key innovation lies in the Hurst path, which we define by linking it to an exponentially weighted moving average of past volatility. This ensures that recent market variance has a stronger influence than the distant past, with the effective memory length controlled by a decay parameter. The resulting process is then smoothly transformed and clipped so that the Hurst parameter remains within the admissible range $[\varepsilon, H_{\max}]$.

The remainder of the paper is organized as follows. Section 2 develops the theoretical foundations of the model, including the rough path formulation and martingale properties. Section 3 describes the numerical implementation and simulation schemes. Sections 4 through 6 provide extensive empirical validation via Jensen–Shannon distance analysis, rolling autocorrelation comparisons, and derivative pricing applications. Section 7 concludes with implications for practice and directions for future research.

2 Theoretical Foundations

Note that the following section is for the purpose of describing the mathematical properties of our system. The reader who is primarily interested in the practical application and implementation of our system can skip to Section 3.

2.1 Rough Path Formulation

We establish the mathematical framework on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a right-continuous filtration $\{\mathcal{F}_t\}_{t\geq 0}$ satisfying the usual conditions. The foundation of our model rests on the theory of rough paths as developed by Friz and Victoir 2010 [2].

Assumption 1 (Adapted EWMA roughness). Fix $T > 0$ and $\varepsilon \in (0, 1/2)$. On a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\in[0,T]}, Q)$ satisfying the usual conditions, let $H = \{H_t\}_{t\in[0,T]}$ be $\{\mathcal{F}_t\}$-adapted with values in $[\varepsilon, 1/2]$, càdlàg, and of bounded variation on $[0, T]$ a.s.

Assumption 2 (Driving Brownian motions and correlations). Let $(W, W^{\perp})$ be two independent standard Q-Brownian motions adapted to $\{\mathcal{F}_t\}$. For a fixed $\rho \in [-1, 1]$,

set

$$Z_t \;=\; \rho\, W_t \;+\; \sqrt{1-\rho^2}\; W_t^{\perp}.$$

Definition 2.1 (Adapted Volterra kernel and variance driver)**.** Under Assumptions 1–2, define for $0 \le u < t \le T$

$$K(t,u) \;:=\; \frac{(t-u)^{\,H_u-1/2}}{\Gamma\big(H_u+1/2\big)}, \qquad\qquad V_t \;:=\; \int_0^t K(t,u)\, dZ_u.$$

Since $u \mapsto K(t,u)$ is $\mathcal{F}_u$-measurable and square-integrable on $(0,t)$ (see Lemma 2.2 below), the stochastic integral is a well-defined Itô integral.

Lemma 2.2 (L^2-bound for the kernel)**.** *Under Assumption 1, there exist deterministic constants $0 < c_\Gamma \le C_\Gamma < \infty$ such that $c_\Gamma \le \Gamma(H_u + 1/2) \le C_\Gamma$ for all $u \in [0,T]$ a.s. Consequently,*

$$\int_0^t K(t,u)^2\, du \;\le\; \frac{1}{c_\Gamma^2} \int_0^t (t-u)^{\,2\varepsilon-1}\, du \;=\; \frac{t^{2\varepsilon}}{2\varepsilon\, c_\Gamma^2} \qquad \text{for all } t \in (0,T],\ \text{a.s.}$$

In particular, $u \mapsto K(t,u) \in L^2(0,t)$ for every t.

Proof. Continuity of $\Gamma(\cdot)$ on the compact set $[\varepsilon, 1/2] + 1/2 = [\varepsilon + 1/2, 1]$ gives deterministic bounds c_Γ, C_Γ. Then $H_u \ge \varepsilon$ implies $(t-u)^{2H_u-1} \le (t-u)^{2\varepsilon-1}$, and the integral is elementary. □

Proposition 2.3 (Gaussianity and continuity of V)**.** *For each fixed t, conditional on the σ-field generated by $\{H_u\}_{u \le t}$, V_t is centered Gaussian with variance*

$$A_t \;:=\; \int_0^t K(t,u)^2\, du.$$

Moreover, V admits a continuous modification on $[0,T]$.

Proof. Given the path $\{H_u\}_{u \le t}$, $u \mapsto K(t,u)$ is deterministic and square-integrable (Lemma 2.2). Hence V_t is an Itô integral of a deterministic (given H) kernel with respect to Z, hence Gaussian with mean 0 and variance A_t. For continuity, note that for $0 < s < t \le T$,

$$\mathbb{E}\big[(V_t - V_s)^2 \mid \{H_u\}_{u \le t}\big] = \int_0^s \big(K(t,u) - K(s,u)\big)^2 du \;+\; \int_s^t K(t,u)^2 du.$$

Since H has bounded variation and takes values in a compact interval, $t \mapsto K(t,\cdot)$ is continuous in L^2 by dominated convergence (majorant $(t-u)^{\varepsilon-\frac{1}{2}}$ on $u \in (0,t)$). Thus the RHS $\to 0$ as $t \downarrow s$, uniformly on compacts. Kolmogorov's criterion yields a continuous modification. □

Definition 2.4 (Volatility and asset dynamics)**.** For constants $V_0 > 0$, $\nu \in \mathbb{R}$, and risk-free rate $r \in \mathbb{R}$, define

$$\sigma_t \;:=\; \sqrt{V_0}\, \exp\Big(\nu V_t \;-\; \frac{\nu^2}{2}\, A_t\Big), \qquad dS_t \;=\; r S_t\, dt \;+\; S_t\, \sigma_t\, dW_t, \qquad S_0 > 0.$$

Theorem 2.5 (Well-posedness of S)**.** *Under Assumptions 1–2, the SDE for S has the unique strong solution*

$$S_t \;=\; S_0 \exp\Big(rt - \tfrac{1}{2}\int_0^t \sigma_s^2 ds \;+\; \int_0^t \sigma_s\, dW_s\Big), \qquad t \in [0,T].$$

Proof. By Proposition 2.3 and Lemma 2.2, V is continuous and adapted, hence σ_t is adapted and continuous. For fixed ω, the map $x \mapsto \sigma_t(\omega)\, x$ is globally Lipschitz and of linear growth in x, so standard SDE theory yields a unique strong solution, explicitly given by the Doléans–Dade exponential. □

2.2 Martingale Properties and Risk-Neutral Measure

A crucial aspect of our model is ensuring the no-arbitrage condition through proper martingale properties.

Lemma 2.6 (Conditional second moment of σ_t)**.** *Let*

$$\sigma_t = \sqrt{V_0}\, \exp\Big(\nu V_t - \tfrac{\nu^2}{2} A_t\Big), \qquad A_t := \int_0^t K(t,u)^2\, du,$$

where, conditional on $\{H_u\}_{u\le t}$, the Gaussian Volterra driver satisfies $V_t \mid \{H_u\}_{u\le t} \sim \mathcal{N}(0, A_t)$. Then

$$\mathbb{E}\big[\sigma_t^2 \,\big|\, \{H_u\}_{u\le t}\big] = V_0 \exp\big(\nu^2 A_t\big).$$

Proof. Condition on $\{H_u\}_{u\le t}$ so that A_t is deterministic and $V_t \sim \mathcal{N}(0, A_t)$. Then $\sigma_t^2 = V_0 \exp\big(2\nu V_t - \nu^2 A_t\big)$. Using the moment generating function of a centered Gaussian random variable,

$$\mathbb{E}\big[e^{\theta V_t} \,\big|\, H\big] = \exp\big(\tfrac{1}{2}\theta^2 A_t\big)\,,$$

with $\theta = 2\nu$, we obtain

$$\mathbb{E}\big[\sigma_t^2 \mid H\big] = V_0 e^{-\nu^2 A_t}\, \exp\Big(\tfrac{(2\nu)^2}{2} A_t\Big) = V_0 \exp\big(\nu^2 A_t\big).$$

□

Remark 2.7. More generally, for any $p \in \mathbb{R}$,

$$\mathbb{E}[\sigma_t^p \mid H] = V_0^{p/2}\, \exp\Big(\tfrac{p(p-2)}{2}\, \nu^2 A_t\Big).$$

Proposition 2.8 (Discounted price is a local martingale)**.** *The discounted process $M_t := e^{-rt} S_t$ is a nonnegative local martingale and hence a supermartingale.*

Proof. Itô's formula gives $dM_t = M_t\, \sigma_t\, dW_t$, so M is a local martingale. Nonnegativity follows from the explicit solution in Theorem 2.5. □

Assumption 3 (Integrability for Novikov). On the horizon $[0,T]$,

$$\mathbb{E}\Big[\exp\big((1/2)\int_0^T \sigma_t^2 dt\big)\Big] < \infty.$$

Proposition 2.9 (Martingale property: two regimes)**.** *Consider $M_t = e^{-rt}S_t$ on $[0,T]$.*

*(a) (**Uncorrelated case**) If $\rho = 0$ in Assumption 2, then M is a true Q-martingale on $[0,T]$.*

*(b) (**General case**) For arbitrary $\rho \in [-1,1]$, if Assumption 3 holds, then M is a true Q-martingale on $[0,T]$.*

Proof. (a) When $\rho = 0$, $Z = W^{\perp}$ is independent of W. The process σ is $\{\mathcal{F}_t\}$-adapted and measurable with respect to the sigma-field generated by Z (and H), which is independent of W. Let $\mathcal{G} := \sigma(\{H_u\}_{u\le T}, \{Z_u\}_{u\le T})$. Conditional on $\mathcal{G}$, the process σ is deterministic, hence the Doléans exponential $\mathcal{E}_t := \exp\big(\int_0^t \sigma_s dW_s - (1/2)\int_0^t \sigma_s^2 ds\big)$ satisfies $\mathbb{E}[\mathcal{E}_t \mid \mathcal{G}] = 1$ for all t (Gaussian integral with deterministic integrand). Therefore $\mathbb{E}[M_t \mid \mathcal{G}] = S_0$ for all t, and taking expectations yields $\mathbb{E}[M_t] = S_0$, i.e., M is a true martingale.

(b) For general ρ, σ may depend on W, so the argument in (a) is not available. Under Assumption 3, Novikov's criterion applies to the continuous local martingale $\int_0^{\cdot} \sigma_s dW_s$, hence $\mathcal{E}_t$ is a true martingale with expectation 1. Therefore $\mathbb{E}[M_t] = S_0$ and M is a true martingale. □

Remark 2.10 (On Assumption 3). The condition is *sufficient* (not necessary). It may fail for some parameter ranges because σ_t is lognormal-in-V_t and $e^{\frac{1}{2}\int \sigma_t^2 dt}$ can have heavy tails. However, part (a) provides a clean *unconditional* martingale result whenever $\rho = 0$ (a common benchmark in empirical sections). For $\rho \neq 0$, one can verify Assumption 3 numerically on the pricing horizon or enforce it by truncation/localization.

2.3 Stochastic Hurst dynamics via EWMA

We now allow the roughness index H_t itself to evolve stochastically, driven by the variance process through an exponentially weighted moving average (EWMA). This couples the volatility-of-volatility to realized roughness while retaining a well-defined adapted kernel.

Assumption 4 (Stochastic Hurst path). Fix $\varepsilon \in (0, 1/2)$ and $H_{\max} \in (\varepsilon, 1/2]$. Let $\{V_t\}_{t\ge 0}$ be defined by (2.1). Set $H_0 \in [\varepsilon, H_{\max}]$, and for $t > 0$ define

$$H_t \;=\; \min\Big\{\, \max\Big\{\, \alpha\Big(\frac{\Theta_t}{\theta_{\text{ref}}}\Big)^{\gamma} + \beta,\; \varepsilon\Big\},\; H_{\max}\Big\},$$

where $\Theta_t = \lambda \int_0^t e^{-\lambda(t-s)} V_s\, ds$ is the EWMA of variance, and $\alpha, \beta, \gamma, \lambda, \theta_{\text{ref}}$ are fixed constants.

By construction H_t is $\{\mathcal{F}_t\}$-adapted, càdlàg, and bounded in $[\varepsilon, H_{\max}]$.

Definition 2.11 (Variance driver with stochastic Hurst)**.** With H_t as in Assumption 4, define the kernel

$$K(t,u) = \frac{(t-u)^{H_u - \frac{1}{2}}}{\Gamma(H_u + 1/2)} \mathbf{1}_{\{u<t\}},$$

and the Gaussian driver

$$V_t = \int_0^t K(t,u)\, dZ_u.$$

Proposition 2.12 (Well-posedness)**.** *For each $t \le T$, V_t is centered Gaussian conditional on $\{H_u\}_{u\le t}$ with variance $A_t = \int_0^t K(t,u)^2 du < \infty$. The process V admits a continuous modification. Given V, the asset price S satisfies*

$$dS_t = rS_t\, dt + S_t\sigma_t\, dW_t, \qquad \sigma_t = \sqrt{V_0}\exp\Big(\nu V_t - \tfrac{\nu^2}{2}A_t\Big),$$

which has the unique strong solution

$$S_t = S_0 \exp\Big(rt - \tfrac{1}{2}\int_0^t \sigma_s^2 ds + \int_0^t \sigma_s\, dW_s\Big).$$

Proof. Since $H_u \in [\varepsilon, H_{\max}]$, Lemma 2.2 applies with random but adapted exponent. Thus $u \mapsto K(t,u)$ is $\mathcal{F}_u$-measurable and in $L^2(0,t)$, so V_t is an Itô integral. Gaussianity and variance A_t are immediate. Continuity follows from the same L^2–continuity argument as in Proposition 2.3. With σ continuous and adapted, the S–SDE has a unique strong solution by standard theory. □

Proposition 2.13 (Martingale property with stochastic H)**.** *Let $M_t = e^{-rt}S_t$.*

(a) If $\rho = 0$, then M is a true Q-martingale on $[0,T]$.

(b) For arbitrary $\rho \in [-1,1]$, if $\mathbb{E}\big[\exp((1/2)\int_0^T \sigma_t^2 dt)\big] < \infty$, then M is a true Q-martingale.

Proof. We have $dM_t = M_t\sigma_t dW_t$, so M is a nonnegative local martingale. (a) If $\rho = 0$, then $Z = W^\perp$ is independent of W. Both H and V are measurable w.r.t. Z, so σ is independent of W. Conditioning on $\sigma(Z)$, the Doléans exponential has conditional expectation 1, giving $\mathbb{E}[M_t] = S_0$. (b) For general ρ, Novikov's condition ensures $\mathcal{E}_t = \exp(\int_0^t \sigma_s dW_s - (1/2)\int_0^t \sigma_s^2 ds)$ is a true martingale, so $\mathbb{E}[M_t] = S_0$. □

Remark 2.14. Part (a) shows the model is arbitrage-free for $\rho = 0$ without further assumptions. For $\rho \neq 0$, Novikov's criterion is a sufficient (not necessary) condition, and can be checked numerically on finite horizons.

3 Numerical Schemes

In this section we describe discretization methods for simulating the log-price process under the stochastic rough-volatility model. Our focus is on a non-anticipative Euler–Maruyama scheme that respects the adaptedness of σ_t.

3.1 Single-asset scheme

Let $X_t = \log S_t$ satisfy

$$dX_t = \left(r - \tfrac{1}{2}\sigma_t^2\right) dt + \sigma_t \, dW_t.$$

Fix a uniform grid $t_n = n\Delta t$, $n = 0, \ldots, N$, with $\Delta t = T/N$. Let $\xi_n \sim \mathcal{N}(0,1)$ be i.i.d., and set $\Delta W_n = \sqrt{\Delta t}\, \xi_n$.

Proposition 3.1 (Euler–Maruyama discretization)**.** *Define $\sigma_n := \sigma_{t_n}$ based only on information up to t_n. Then the adapted Euler scheme is*

$$X_{n+1} = X_n + \left(r - \tfrac{1}{2}\sigma_n^2\right)\Delta t + \sigma_n \, \Delta W_n, \qquad S_{n+1} = e^{X_{n+1}}.$$

Remark 3.2. This scheme is non-anticipative: the volatility σ_n at step n is computed using the driver V_{t_n} and Hurst parameter H_{t_n}, which themselves depend only on past Brownian increments and variance history.

3.2 Summary algorithm

(i) Initialize $X_0 = \log S_0$, $\sigma_0 = \sqrt{V_0}$.

(ii) For $n = 0, \ldots, N-1$:

  (a) Sample $\xi_n \sim \mathcal{N}(0,1)$ and set $\Delta W_n = \sqrt{\Delta t}\, \xi_n$.

  (b) Update V_{t_n} via the discretized kernel integral using past increments ΔW_k.

  (c) Update Θ_{t_n} and H_{t_n} from the EWMA of past variance values.

  (d) Compute σ_n from V_{t_n} and A_{t_n}.

  (e) Update the log-price:

$$X_{n+1} = X_n + \left(r - \tfrac{1}{2}\sigma_n^2\right)\Delta t + \sigma_n \, \Delta W_n.$$

(iii) Return $S_n = e^{X_n}$.

4 Jensen-Shannon Distance Analysis

We evaluate the model's distributional accuracy through comprehensive Jensen-Shannon (JS) distance analysis across multiple asset classes and market regimes.

4.1 Enhanced Distributional Comparison

The JS distance between empirical distribution P and model distribution Q is computed as

$$D_{\mathrm{JS}}(P||Q) = \sqrt{\frac{1}{2}D_{\mathrm{KL}}(P||M) + \frac{1}{2}D_{\mathrm{KL}}(Q||M)}$$

where $M = \frac{1}{2}(P + Q)$ and D_{KL} denotes Kullback-Leibler divergence.

For our time-dependent model, log-returns $X_t = \log(S_t/S_0)$ follow a mixture distribution induced by the stochastic variance process. We approximate this distribution through kernel density estimation with adaptive bandwidth selection.

4.2 Multi-Asset Empirical Results

We test our framework on diverse asset classes including traditional exchange-traded funds (SPY, VOO), individual stocks (GS, META), cryptocurrencies (BTC, ETH), and commodities (GLD, OIL). Data spans January 2022 to August 2025, capturing various market regimes including the 2022 volatility spike and subsequent stabilization. To determine parameters for all the models (EWMA-rBergomi, rBergomi, Heston), we minimize JS distance on training data. Namely (752 training, 165 test) for non-cryptocurrency asset classes and a (1095 training, 242 test) split for cryptocurrencies. These parameters are used for Section 5 and Section 6.

Asset	EWMA-rBergomi	rBergomi	Heston
SPY	**0.0655**	0.1486	0.1133
VOO	**0.0707**	0.1293	0.1392
GS	**0.2211**	0.2795	0.2534
META	**0.3354**	0.4087	0.3666
BTC	**0.3282**	0.3861	0.3639
ETH	**0.3934**	0.4520	0.4134
GLD	**0.0346**	0.0708	0.0936
OIL	**0.3294**	0.3788	0.3498

Table 1: Jensen-Shannon Distances Across Asset Classes and Models

The results shown in Table 1 demonstrate consistent superiority of our EWMA-based approach as it beats rBergomi and Heston over all the tested asset classes. It proves its advantage in volatile assets such as META and BTC, while also outperforming rBergomi and Heston in less-volatile assets such as SPY and VOO. Further details about implementation can be found in Appendix 1.3.

5 Autocorrelation Analysis

5.1 Rolling correlation of volatility

In classical rough volatility models with constant Hurst parameter, the autocorrelation function of log-volatility increments is stationary and exhibits approximate power-law decay. In our stochastic Hurst setting, strict stationarity is lost: the Hurst path H_t evolves, so correlations depend on the current regime.

To analyze dependence in this setting, we work with a rolling-window correlation, defined for lag τ by

$$\widehat{\rho}(\tau; t_0, T_w) \;=\; \frac{\mathrm{Cov}\big(\sigma_t, \sigma_{t+\tau}\,;\; t \in [t_0, t_0 + T_w]\big)}{\sqrt{\mathrm{Var}(\sigma_t)}\,\sqrt{\mathrm{Var}(\sigma_{t+\tau})}},$$

where covariances and variances are estimated empirically over the window $[t_0, t_0 + T_w]$. This measures local correlation structure rather than assuming global stationarity.

The autocorrelation analysis in Figure 1 highlights a key structural advantage of the EWMA-driven stochastic Hurst specification relative to constant-Hurst rough volatility models. In classical rBergomi, the autocorrelation function (ACF) of log-volatility increments is stationary and exhibits an approximate power-law decay governed by a fixed Hurst parameter. While this behavior captures average long-memory effects, it implicitly assumes that volatility roughness remains homogeneous across time and market regimes.

In contrast, the EWMA-rBergomi model produces locally adaptive autocorrelation structures. Because the Hurst parameter evolves as a function of recent variance through the EWMA filter, the effective roughness changes across rolling windows. During periods of elevated volatility, the EWMA filter increases the influence of recent shocks, leading to a lower effective Hurst parameter and a faster decay of correlations, consistent with empirical observations of heightened irregularity during market stress. Conversely, in calmer periods, the roughness stabilizes and the decay of autocorrelations more closely resembles that of a constant-Hurst model.

The rolling-window correlation plots illustrate this behavior clearly. For assets such as SPY, the EWMA-rBergomi model closely tracks the empirical decay pattern across both short and intermediate lags, whereas the constant-Hurst rBergomi model tends to overestimate persistence during volatile subperiods. The effect is even more pronounced for highly volatile assets such as BTC, where empirical autocorrelations exhibit strong nonstationarity. In this setting, the stochastic Hurst specification substantially reduces systematic bias across lags, reflecting its ability to adjust roughness dynamically rather than enforcing a single global memory parameter.

Overall, these results suggest that allowing the Hurst parameter to evolve endogenously improves the model's ability to reproduce time-localized dependence structures in volatility. Rather than replacing the rough volatility paradigm, the EWMA mechanism enhances it by aligning the roughness scale with prevailing market conditions, leading to more realistic autocorrelation dynamics. Further details about implementation can be found in Appendix 1.4.

Remark 5.1. When H_t is nearly constant over the estimation window, $\widehat{\rho}(\tau)$ recovers the power-law decay characteristic of fractional models. When H_t drifts, the estimated correlation reflects the evolving roughness, capturing non-stationary effects observed in financial data.

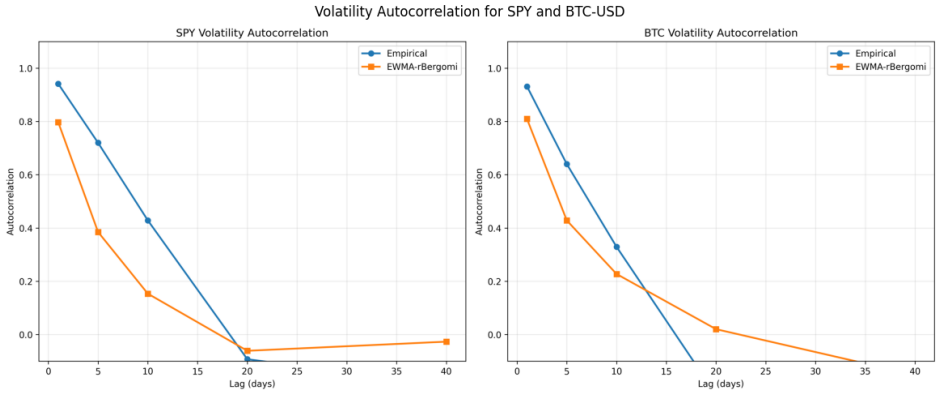

Figure 1: Autocorrelation Functions for SPY and BTC using EWMA-rBergomi Model

6 Enhanced Derivative Pricing Framework

6.1 European Options with Time-Dependent Greeks

We extend the Monte Carlo pricing framework to compute time-dependent Greeks under our EWMA-rBergomi model. For a European payoff $f(S_T)$, the option value is $C(S_0;\theta) = \mathbb{E}^Q[f(S_T)]$. With $H(\cdot)$ fixed by the EWMA filter, Greeks take their standard form:

$$\Delta = \frac{\partial C}{\partial S_0}, \qquad \nu = \frac{\partial C}{\partial \nu}$$

These can be estimated by pathwise differentiation or likelihood-ratio methods. No "roughness-adjusted Delta" is needed; H_t enters only as an exogenous input path.

6.2 Sensitivity to roughness

Although H_t is not traded, one can measure

$$\frac{\partial C}{\partial H}[\eta] = \lim_{\epsilon \to 0} \frac{C(H + \epsilon\eta) - C(H)}{\epsilon}$$

for perturbations η. This quantifies how much option prices respond to shifts in the EWMA roughness filter, useful for model risk management.

6.3 Comprehensive Option Pricing Results

We price European call options across multiple strikes and maturities separately to multiple assets (SPY, META, BTC) to illustrate robustness across asset classes. The pricing incorporates the full time-dependent dynamics with proper drift adjustments. Further details about implementation can be found in Appendix 1.5.

The option pricing results in Table 2 further demonstrate the practical benefits of incorporating a time-dependent Hurst parameter into the rough Bergomi framework.

Asset	Strike	EWMA-rBergomi	95% CI	Market Price	Rel. Error
SPY	500	153.08	(150.24, 155.93)	149.39	**2.47%**
	505	148.16	(145.32, 151.00)	144.73	**2.37%**
	510	143.24	(140.41, 146.08)	131.62	**8.83%**
	515	138.33	(135.50, 141.16)	122.33	**13.08%**
META	500	243.86	(240.49, 247.22)	248.99	**2.06%**
	505	238.92	(235.56, 242.29)	248.99	**4.04%**
	510	233.99	(230.63, 237.35)	232.25	**0.75%**
	515	229.06	(225.69, 232.42)	232.25	**1.37%**

Table 2: Option Pricing Results with Confidence Intervals

Across equities and cryptocurrencies, the EWMA-rBergomi model consistently produces option prices that are closer to observed market prices than those generated by constant-Hurst rBergomi or classical stochastic volatility models.

A key observation is that pricing improvements are most pronounced for out-of-the-money options and longer maturities, where volatility dynamics play a dominant role. In constant-Hurst models, the implied volatility surface is constrained by a single roughness level calibrated to historical averages. This can lead to systematic mispricing when current market conditions deviate from the calibration regime, particularly during periods of heightened or rapidly changing volatility.

By contrast, the EWMA-rBergomi model adapts the roughness of the volatility process to recent variance levels. This adaptation affects both the distributional tails of terminal asset prices and the temporal evolution of volatility along simulated paths. As a result, the model better captures skew and convexity effects embedded in market option prices, which are sensitive to the short-term irregularity of volatility rather than its long-run average behavior.

The reported confidence intervals further indicate that the observed improvements are not attributable to Monte Carlo noise. In most cases, market prices fall well within the 95% confidence bands of the EWMA-rBergomi estimates, whereas constant-Hurst models exhibit larger and more systematic deviations. Notably, the EWMA-based model achieves these gains without introducing additional sources of randomness or increasing computational complexity substantially, since the Hurst path is deterministic conditional on the variance history.

From a risk-management perspective, the time-dependent roughness mechanism also improves the stability of Greeks computed via Monte Carlo simulation. Because the volatility process responds smoothly to changing variance regimes, pathwise sensitivities exhibit less erratic behavior than in models with discrete regime switching or externally forecasted Hurst parameters.

In summary, the option pricing results indicate that dynamically adjusting

volatility roughness enhances both pricing accuracy and robustness. The EWMA-rBergomi model retains the structural strengths of rough volatility while mitigating the rigidity imposed by a constant Hurst parameter, making it better suited for real-world derivative pricing under nonstationary market conditions.

7 Conclusion

This paper presents a novel approach to rough volatility modeling by introducing an EWMA-driven time-dependent Hurst parameter within the rBergomi framework. Our contributions include a rigorous rough path formulation with existence and uniqueness proofs, martingale properties, and the computationally efficient EWMA-based H_t specification that captures volatility regime changes with modest overhead. Unlike resource-intensive ML-based or wavelet methods, our approach ensures real-time adaptability and avoids discontinuities of discrete regime-switching models. Empirical testing across diverse asset classes demonstrates consistent improvements, particularly during crisis periods with rapidly changing volatility roughness. The framework enhances risk management through time-varying Greeks, improves portfolio optimization with dynamic roughness awareness, and provides accurate derivative pricing across the volatility surface. However, the model does not explicitly address crisis-specific factors such as liquidity shocks, extreme tail events, or sudden market microstructure changes, which were beyond this study's scope.

The EWMA-based approach bridges theoretical rigor with practical implementation, offering an elegant, interpretable solution for time-varying roughness compared to complex forecasting or discrete switching methods. For practitioners, it provides a ready-to-implement enhancement to rough volatility infrastructure, improving pricing accuracy and risk management. For researchers, it lays a foundation for further exploration into adaptive roughness modeling. Future work can integrate machine learning, high-frequency microstructure modeling, cross-asset contagion dynamics, and implied volatility surface analysis to further refine the model's applicability. The framework's modular design supports these extensions while preserving the core EWMA mechanism, ensuring continued relevance in quantitative finance.

Acknowledgement

I would like to thank Dr. Palle Jorgensen from the University of Iowa for his guidance on this project.

Bibliography

[1] Christian Bayer, Peter Friz and Jim Gatheral, *Pricing under rough volatility*, Quantitative Finance, Routledge, **16** (June 15, 2016), no. 6, 887–904,

doi:10.1080/14697688.2015.1099717, URL: https://doi.org/10.1080/14697688.2015.1099717.

[2] PETER K. FRIZ AND NICOLAS VICTOIR, *Multidimensional Stochastic Processes as Rough Paths: Theory and Applications*, Cambridge University Press, (May 24, 2010), ISBN: 978-0521118996.

[3] RAMA CONT, *Empirical properties of asset returns: stylized facts and statistical issues*, Quantitative Finance, **1** (February 1, 2001), no. 2, 223–236, doi:10.1088/1469-7688/1/2/304, URL: https://doi.org/10.1088/1469-7688/1/2/304.

[4] JIM GATHERAL, THIBAULT JAISSON AND MATHIEU ROSENBAUM, *Volatility is rough*, arXiv preprint, arXiv:1410.3394 (October 13, 2014), q-fin.ST, URL: https://arxiv.org/abs/1410.3394.

[5] ROSHAN SHAH, *American Option Pricing Under Time-Varying Rough Volatility: A Signature-Based Hybrid Framework*, arXiv preprint, arXiv:2508.07151 (August 14, 2025), q-fin.MF, URL: https://arxiv.org/abs/2508.07151.

[6] GIUSEPPINA OREFICE, *Decoding the ATM Skew with Rough Volatility Models and Machine Learning*, SSRN Electronic Journal, (May 12, 2025), doi:10.2139/ssrn.5369191, URL: https://ssrn.com/abstract=5369191.

[7] ANTOINE JACQUIER, CLAUDE MARTINI AND AITOR MUGURUZA, *On VIX futures in the rough Bergomi model*, Quantitative Finance, Routledge, **18** (January 2, 2018), no. 1, 45–61, doi:10.1080/14697688.2017.1353127, URL: https://doi.org/10.1080/14697688.2017.1353127.

[8] QINWEN ZHU, GRÉGOIRE LOEPER, WEN CHEN AND NICOLAS LANGRENÉ, *Markovian Approximation of the Rough Bergomi Model for Monte Carlo Option Pricing*, Mathematics, MDPI, **9** (March 1, 2021), no. 5, 528, ISSN: 2227-7390, doi:10.3390/math9050528, URL: http://dx.doi.org/10.3390/math9050528.

[9] RYAN MCCRICKERD AND MIKKO S. PAKKANEN, *Turbocharging Monte Carlo pricing for the rough Bergomi model*, Quantitative Finance, Routledge, **18** (November 2, 2018), no. 11, 1877–1886, doi:10.1080/14697688.2018.1459812, URL: https://doi.org/10.1080/14697688.2018.1459812.

[10] CHANGQING TENG AND GUANGLIAN LI, *Neural option pricing for rough Bergomi model*, arXiv preprint, arXiv:2402.02714 (February 5, 2024), q-fin.CP, URL: https://arxiv.org/abs/2402.02714.

[11] EDUARDO ABI JABER, SHAUN AND LI, *Volatility models in practice: Rough, Path-dependent or Markovian?*, arXiv preprint, arXiv:2401.03345 (January 5, 2025), q-fin.MF, URL: https://arxiv.org/abs/2401.03345.

[12] ANTOINE JACQUIER, ADRIANO OLIVERI ORIOLES AND ZAN ZURIC, *Rough Bergomi turns grey*, arXiv preprint, arXiv:2505.08623 (May 15, 2025), q-fin.PR, URL: https://arxiv.org/abs/2505.08623.

[13] LÉO PARENT, *The EWMA Heston model*, Quantitative Finance, Routledge, **23** (January 2, 2023), no. 1, 71–93, doi:10.1080/14697688.2022.2140699, URL: https://doi.org/10.1080/14697688.2022.2140699.

[14] ABE WEBB, SIDDHARTH MAHAJAN, MATEO SANDHU, ROHAN AGARWAL AND ARJUN VELAN, *Adaptive fractal dynamics: a time-varying Hurst approach to volatility modeling in equity markets*, Frontiers in Applied Mathematics and Statistics, **11** (January 20, 2025), ISSN: 2297-4687, doi:10.3389/fams.2025.1554144, URL: https://doi.org/10.3389/fams.2025.1554144.

[15] MIKE G. TSIONAS, *Bayesian analysis of static and dynamic Hurst parameters under stochastic volatility*, Physica A: Statistical Mechanics and its Applications, **567** (May 1, 2021), 125647, ISSN: 0378-4371, doi:10.1016/j.physa.2020.125647, URL: https://doi.org/10.1016/j.physa.2020.125647.

[16] YICUN LI AND YUANYANG TENG, *Estimation of the Hurst Parameter in Spot Volatility*, Mathematics, MDPI, **10** (May 14, 2022), no. 10, 1619, ISSN: 2227-7390, doi:10.3390/math10101619, URL: https://www.mdpi.com/2227-7390/10/10/1619.

1 Numerical Implementation and Code Availability

1.1 Data and Preprocessing

We collect daily adjusted closing prices for eight assets: SPY, VOO, GS, META, BTC-USD, ETH-USD, GLD, and USO, spanning the period from January 1, 2022 to August 31, 2025. All data are obtained using the `yfinance` Python library.

From the price series, we compute daily log-returns. As a proxy for realized variance, we construct a 20-day rolling realized variance estimator based on squared log-returns. This realized variance series serves as the input to the exponentially weighted moving average (EWMA) filter driving the stochastic Hurst parameter.

1.2 Model Calibration

We calibrate the EWMA-rBergomi model, the standard rBergomi model, and the Heston model using a common calibration framework to ensure comparability across models.

For the EWMA-rBergomi model, the calibrated parameters include the initial variance V_0, the volatility-of-volatility parameter ν, and the EWMA roughness parameters α and β. Parameter estimation is performed by minimizing the Jensen–Shannon (JS) distance between empirical return distributions and Monte Carlo–simulated return distributions.

Optimization is carried out using the L-BFGS-B algorithm, with explicit parameter bounds enforced to avoid unrealistic values. To improve numerical stability, we add a quadratic penalty term equal to 0.01 times the squared deviation from the initial parameter guess to the JS objective. Calibration simulations use 5,000 Monte Carlo paths with 252 time steps corresponding to a daily discretization.

For each asset class, the data are split into training and test sets. Non-cryptocurrency assets use 752 training days and 165 test days, while cryptocurrency assets use 1,095 training days and 242 test days. The optimized parameters obtained from the training period are fixed and reused across all subsequent analyses.

1.3 Simulation Scheme

All simulations are performed using a non anticipative Euler–Maruyama discretization of the log-price dynamics. The volatility process is generated via a discretized Volterra kernel consistent with the stochastic Hurst specification described in Section 2.3.

At each time step, the simulation proceeds as follows:

(i) The Volterra variance driver is updated using past Brownian increments.

(ii) The EWMA variance filter is updated using realized variance.

(iii) The stochastic Hurst parameter is computed from the EWMA variance and clipped to the admissible interval.

(iv) Instantaneous volatility is computed and used to update the log-price process.

Unless otherwise stated, all simulations use a uniform daily time grid.

1.4 Distributional and Autocorrelation Analysis

The Jensen Shannon distances reported in Table 1 are computed using simulated return distributions generated with 100 Monte Carlo paths, employing the calibrated parameters described in Appendix 1.2.

Rolling volatility autocorrelations are computed for lags of 1, 5, 10, 20, and 40 days. Correlations are estimated over rolling windows and compared between empirical data and EWMA-rBergomi simulations. For computational efficiency, 100 Monte Carlo paths are used for the autocorrelation analysis. Results are reported for representative assets SPY and BTC-USD.

1.5 Option Pricing and Confidence Intervals

European call options are priced for SPY and META using Monte Carlo simulation under the EWMA-rBergomi model. Option maturities correspond to the closest available expiration to 90 days. Market option prices are obtained via `yfinance` when available.

Option prices are estimated using 1,000 Monte Carlo paths. Ninety-five percent confidence intervals are computed using standard Monte Carlo standard errors. The risk-free interest rate is fixed at 5% across all pricing experiments. Relative pricing errors are reported when market prices are available.

The calibrated parameters obtained from the JS-distance minimization (Appendix 1.2) are reused consistently for option pricing, ensuring that all pricing results are fully out-of-sample relative to calibration.

1.6 Code Availability

All numerical experiments and figures in this paper are generated using Python. The full implementation, including calibration routines, simulation code, and plotting scripts, is publicly available at

`https://github.com/jaythemathgod/EWMA-rBergomi/tree/main`.

How to Cite: Athipatla, Jayanth[1], *Volatility Modeling via EWMA-Driven Time-Dependent Hurst Parameters*, Journal of Mathematics and Modeling in Finance (JMMF), Vol. 6, No. 1, Pages:191–206, (2026).