Variance-Exact Simulation of Rough Volatility: A Comparison of Hybrid, Davies-Harte, and L^2-Exact Riemann-Liouville Schemes

M.A. van den Berg
simu.ai
thijs@simu.ai

November 6, 2025

Abstract

We compare three principal numerical schemes for simulating fractional Gaussian drivers used in rough volatility models: the classical hybrid method of Bayer et al. [2016], the exact circulant embedding method of Davies and Harte [1987], and a newly introduced L^2-exact Riemann-Liouville (RL) convolution scheme. While the hybrid method remains dominant in finance due to its causal structure, we show that it systematically underestimates the variance for small Hurst exponents ($H \lesssim 0.1$). The proposed L^2-exact modification preserves causality while enforcing variance-matching in the near field, achieving accuracy comparable to Davies-Harte but with the correct RL integral form required in models such as rough Bergomi and rough Heston.

1 Introduction

Rough volatility models such as the rough Bergomi model [Bayer et al., 2016] and rough Heston [Gatheral et al., 2018] rely on Gaussian processes with long-memory kernels of the form

$$Y_t = \sqrt{2H} \int_0^t (t-s)^{H-\frac{1}{2}} \, dW_s, \tag{1}$$

with $H < \frac{1}{2}$. Efficient simulation of Y_t on a discrete time grid is therefore essential for Monte Carlo methods in rough volatility.

Three families of algorithms dominate the literature:

1. **Hybrid scheme (BLP)**: Introduced by Bayer et al. [2016], combines exact bin integration near zero with point sampling in the tail.
2. **Davies-Harte (DH)**: Uses a circulant embedding of the covariance matrix to generate exact discrete fractional Brownian motion (fBM) samples [Davies and Harte, 1987].
3. **L^2-exact RL scheme (this work)**: Replaces near-field bin averages with variance-exact weights that preserve the L^2 energy of each kernel segment.

2 Comparison of Methods

Table 1 summarizes the essential properties.

The hybrid method is widely used in finance because it is simple, causal, and compatible with FFT convolution. However, for small H the near-field bin averaging introduces a systematic

Property	Hybrid (BLP)	Davies-Harte	L^2-Exact RL
Target process	RL integral	fBM	RL integral
Causality	Yes	No	Yes
Covariance exactness	Approx.	Exact	Local L^2 exact
Variance bias (small H)	High	None	None
Streaming simulation	Possible	Impossible	Possible
FFT complexity	$O(N \log N)$	$O(N \log N)$	$O(N \log N)$

Table 1: Qualitative comparison of rough-kernel simulation schemes.

variance deficit: the energy of the kernel in the first few bins is underestimated. Davies-Harte, in contrast, produces an *exact* covariance match for fBM but is non-causal and therefore unsuitable for stochastic integration in SDEs such as (1).

The L^2-exact scheme proposed here keeps the hybrid structure but replaces the near-field weights h_j with

$$h_j = \sqrt{\frac{2H}{\Delta t}} \left[(j\Delta t)^{2H} - ((j-1)\Delta t)^{2H} \right]^{1/2}, \qquad j = 1, \ldots, \kappa, \tag{2}$$

which preserves the variance $\mathbb{E}[|h_j \Delta W_j|^2]$ of each bin exactly. The tail region $(j > \kappa)$ still uses point evaluations $h_j = \sqrt{2H}\,(j\Delta t)^{H-\frac{1}{2}}$ for numerical stability and causality.

3 Empirical Results

Figure 1 compares the empirical standard deviation of the simulated Y_t for multiple H with the theoretical t^H scaling law. The L^2-exact version matches the continuum variance within machine precision, while the standard hybrid scheme shows persistent underestimation for small H even as the grid is refined.

4 Discussion

The Davies–Harte method remains the gold standard for simulating fractional Brownian motion (fBM) exactly on a discrete grid: the finite-dimensional covariance structure is reproduced without error. However, its construction relies on a global circulant embedding of the covariance matrix and therefore generates the entire trajectory at once. Consequently, the process is *non-causal* and has no well-defined driving Brownian motion W_t such that $B_H(t) = \int_0^t K_H(t-s)\,dW_s$. This loss of coupling to a single underlying Wiener path is harmless when one studies marginal laws or variance scaling, but it becomes problematic in models such as rough Bergomi, where the same Brownian motion also drives the asset price:

$$dX_t = \sqrt{V_t}\,dW_t^{(S)}, \qquad V_t = \xi_0 \exp\Big(\eta\, Y_t - \tfrac{1}{2}\eta^2 t^{2H}\Big), \quad Y_t = \sqrt{2H} \int_0^t K_H(t-s)\,dW_s^{(V)}. \tag{3}$$

Here, the volatility factor Y_t and the price driver $W_t^{(S)}$ may be correlated (via ρ), which requires a *causal* construction of Y_t in terms of the same underlying Brownian paths. The Davies–Harte samples cannot be written in this form; they correspond to an independent fBM that shares the same marginals but not the filtration of W_t.

The classical hybrid scheme of Bayer et al. [2016] and our proposed L^2-exact variant both preserve this causal coupling. They compute the Riemann-Liouville convolution directly in discrete time, producing Y_t as an explicit functional of the past Brownian increments. The trade-off is that the hybrid scheme only approximates the true kernel, whereas the L^2-exact

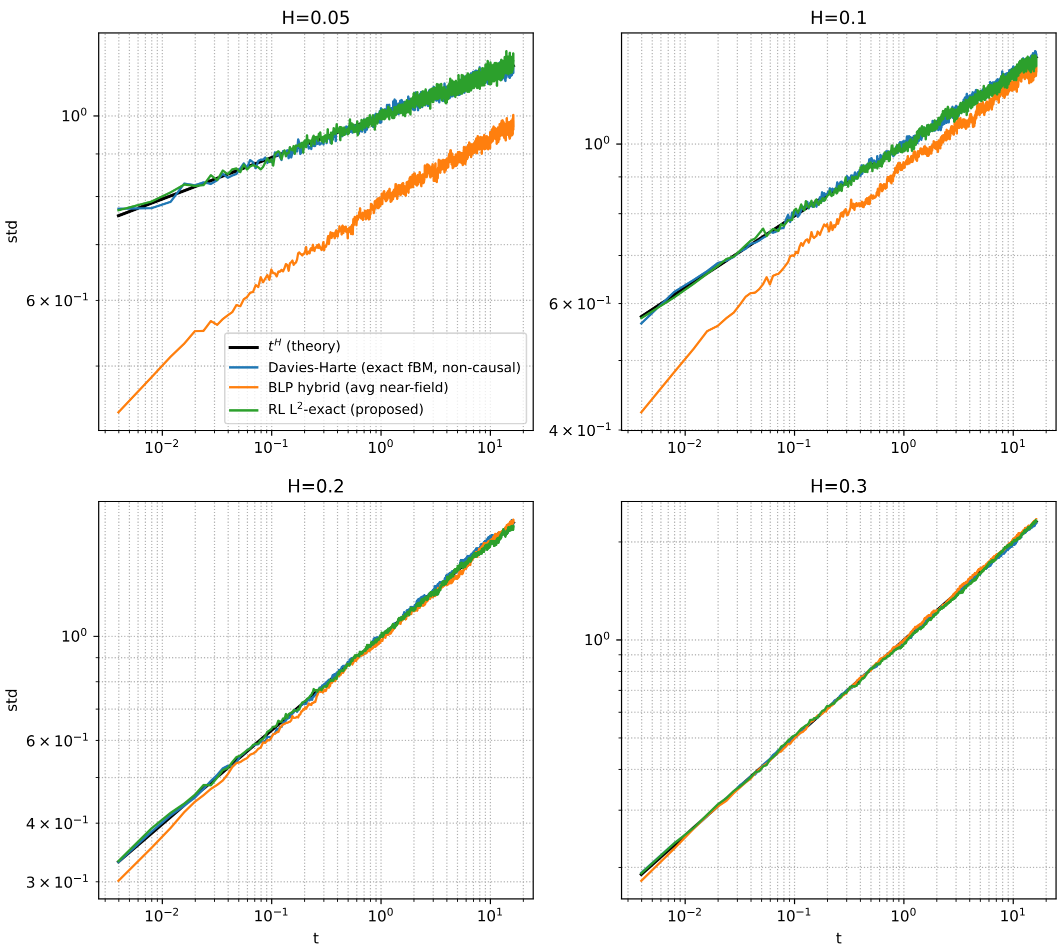

Figure 1: Empirical $\mathrm{std}(Y_t)$ for Davies-Harte (blue), Hybrid (orange), and L^2-Exact (green) schemes across Hurst exponents $H \in \{0.05, 0.1, 0.2, 0.3\}$.

modification enforces variance consistency in the near field and therefore eliminates the local underdispersion that arises for small Hurst exponents. The resulting process retains the correct causal dependence structure while achieving the same variance accuracy as Davies-Harte on the grid.

5 Conclusion

The L^2-exact hybrid scheme provides a variance-consistent, causal, and computationally efficient alternative to both the standard hybrid and the Davies-Harte methods. It aligns with the Riemann-Liouville formulation of rough volatility and avoids the under-dispersion seen in existing implementations for small H.

References

C. Bayer, P. K. Friz, and J. Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6):887-904, 2016.

R. B. Davies and D. S. Harte. Tests for Hurst effect. *Biometrika*, 74(1):95-101, 1987.

J. Gatheral, T. Jaisson, and M. Rosenbaum. Volatility is rough. *Quantitative Finance*, 18(6):933-949, 2018.

K. Diethelm. *The Analysis of Fractional Differential Equations*. Springer, 2010.

C. Li and W. Deng. A review on numerical methods for fractional stochastic differential equations. *Applied Mathematical Modelling*, 77:132-155, 2020.