Replication of “Multivariate Fractional Brownian Motion for Realized Volatility Forecasting”

Bibinger, Yu & Zhang (2025)

Yassine M., Milos G., Issam F. & Yves-Marie S.

May 28, 2026

Abstract

This paper replicates the main theoretical and empirical results of Bibinger et al. (2025), who introduce a multivariate fractional Brownian motion (mfBm) framework for modelling and forecasting realized volatility. We reproduce all seven figures, all eleven main tables, and the Appendix B panel of 20 DJ30 stocks using a Python implementation of the BYZ estimators, the optimal forecasting rule, and the Monte Carlo and empirical studies. All code is available at `https://github.com/yessinemx/Mfbm_realized_vol`.

Contents

1 Introduction

Bibinger et al. (2025) propose a time-reversible multivariate fractional Brownian motion as a continuous-time model for log realized volatility. The model generalises the univariate fBm of Gatheral et al. (2018) to the multivariate case and provides a closed-form optimal predictor.

The rest of this paper is organized as follows. Section 2 recalls the model. Section 3 describes the estimators. Section 4 presents the optimal forecasting rule. Section 5 reproduces the Monte Carlo study. Section 6 presents the empirical results. Section 7 concludes.

2 The Model

2.1 Univariate fractional Brownian motion

A fractional Brownian motion $B^H = (B^H_t)_{t\geq 0}$ with Hurst exponent $H \in (0,1)$ is a centred Gaussian process with covariance

$$\mathrm{Cov}(B^H_s, B^H_t) = \tfrac{1}{2}\Big(|s|^{2H} + |t|^{2H} - |s-t|^{2H}\Big). \tag{1}$$

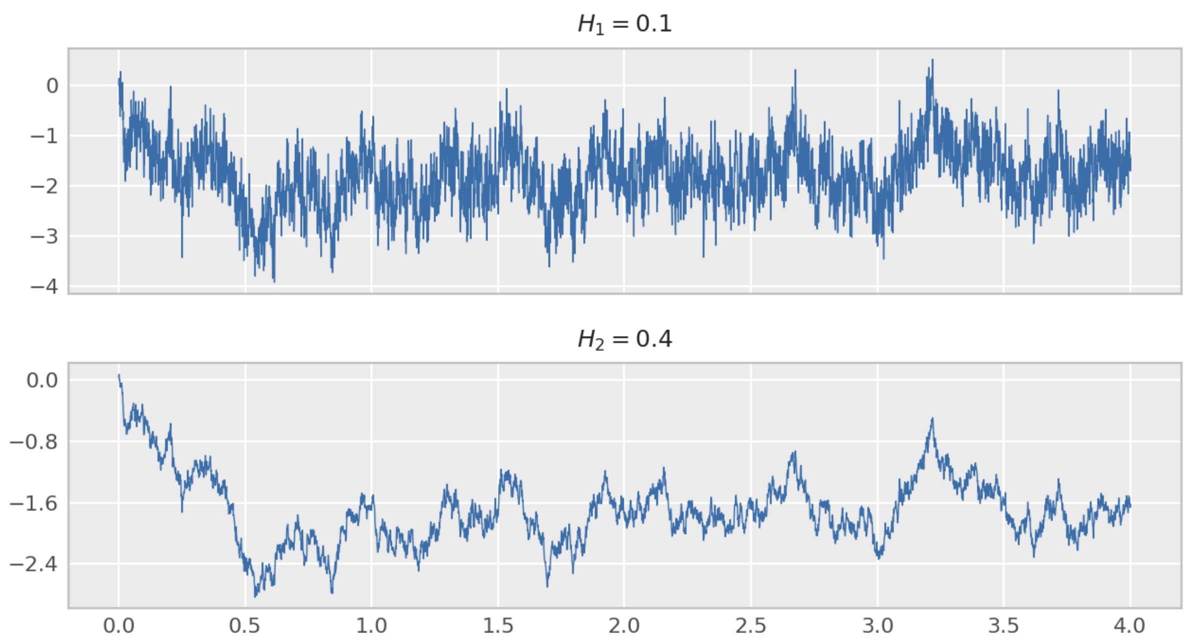

Figure 1: Simulated sample paths of the two components of a bivariate time-reversible mfBm with $H_1 = 0.1$ (top panel, rough) and $H_2 = 0.4$ (bottom panel, smoother), $\rho = 0.4$, $\eta = 0$, $\sigma_1 = \sigma_2 = 1$, $n = 1\,000$ steps. The top trajectory is far more irregular than the bottom one, illustrating the effect of the Hurst exponent on path regularity. Figure 1 of Bibinger et al. (2025).

2.2 Bivariate time-reversible mfBm

Define the auxiliary kernel

$$w(t,h,H) = \tfrac{1}{2}\Big(|t+h|^{2H} + |t|^{2H} - |h|^{2H}\Big), \tag{2}$$

so that $\mathrm{Cov}(B^H_{t+h}, B^H_t) = \sigma^2 w(t,h,H)$. The general cross-covariance of a bivariate mfBm with parameters $(H_1, H_2, \sigma_1, \sigma_2, \rho, \eta)$ is (Eq. (4) of Bibinger et al. (2025), valid for $H_1 + H_2 \neq 1$):

$$\mathrm{Cov}\Big(B_s^{(1)}, B_t^{(2)}\Big) = \frac{\sigma_1\sigma_2}{2}\Big[(\rho+\eta\,\mathrm{sgn}(s))\,|s|^{H_1+H_2} + (\rho-\eta\,\mathrm{sgn}(t))\,|t|^{H_1+H_2} - (\rho-\eta\,\mathrm{sgn}(t-s))\,|t-s|^{H_1+H_2}\Big]. \tag{3}$$

For the *time-reversible* case $\eta = 0$ (Eq. (5) of Bibinger et al. (2025)), this simplifies to

$$\mathrm{Cov}\Big(B_s^{(1)}, B_t^{(2)}\Big)\Big|_{\eta=0} = \rho\,\sigma_1\sigma_2\, w\Big(t,\ s-t,\ \tfrac{H_1+H_2}{2}\Big). \tag{4}$$

2.3 Admissible parameter space

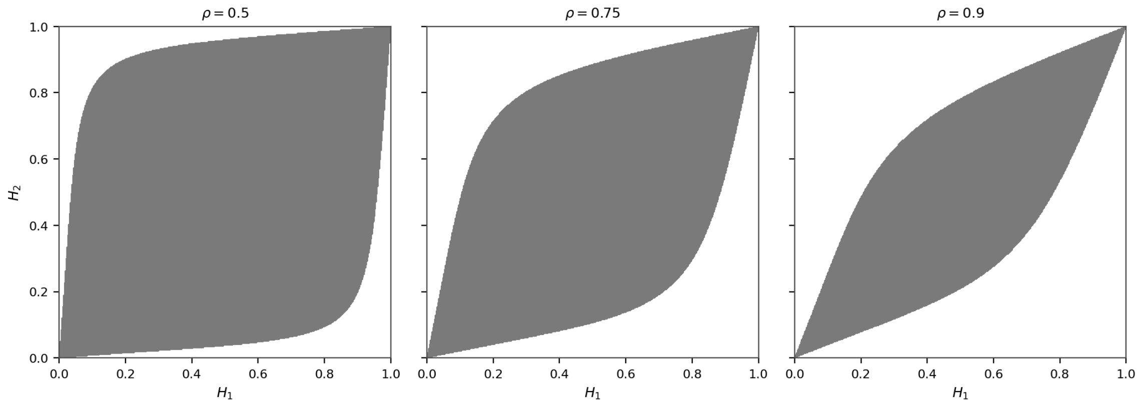

Figure 2: Admissible region of (H_1, H_2) satisfying $|\rho| \leq \rho_{\max}(H_1, H_2)$ for $\rho \in \{0.5,\ 0.75,\ 0.9\}$ (grey = admissible, white = inadmissible). As $|\rho|$ increases, the constraint tightens and pairs with very different Hurst exponents become inadmissible. Figure 3 of Bibinger et al. (2025). Replication via `make_figures.py`.

The maximum admissible correlation $\rho_{\max}(H_1, H_2)$ is illustrated in Figure 3.

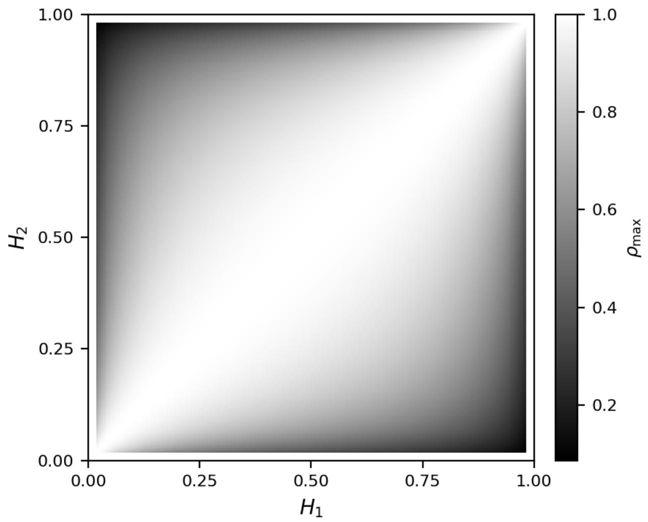

Figure 3: Heatmap of the maximum admissible correlation $\rho_{\max}(H_1, H_2)$ over $(H_1, H_2) \in [0,1]^2$. Lighter shading indicates a larger admissible range; $\rho_{\max} = 1$ on the diagonal $H_1 = H_2$ and decreases as $|H_1 - H_2|$ grows. Figure 2 of Bibinger et al. (2025). Replication via `make_figures.py`.

3 Estimators

3.1 Hurst exponent

The BYZ estimator of H_1 based on observations at times $\delta, 2\delta, \ldots, n\delta$ is (Eq. (6) of Bibinger et al. (2025)):

$$\hat{H}_1 = \frac{1}{2\log 2} \log \frac{\sum_{k=1}^{n/2} (B^{(1)}_{2k\delta} - B^{(1)}_{(2k-2)\delta})^2}{\sum_{k=1}^{n} (B^{(1)}_{k\delta} - B^{(1)}_{(k-1)\delta})^2}. \tag{5}$$

The asymptotic variance of $\hat{H}_1$ is given in Theorem 3.1 of Bibinger et al. (2025).

3.2 Scale, correlation and asymmetry

The variance parameter is estimated by (Eq. (7) of Bibinger et al. (2025)):

$$\hat{\sigma}_j^2 = \frac{\sum_{k=1}^{n} (\Delta_1 B^{(j)}_k)^2}{n\, \delta^{2\hat{H}_j}}. \tag{6}$$

The contemporaneous correlation is estimated by (Eq. (8a) of Bibinger et al. (2025)):

$$\hat{\rho} = \frac{\sum_{k=1}^{n} \Delta_1 B_k^{(1)} \, \Delta_1 B_k^{(2)}}{\sqrt{\sum_{k=1}^{n} (\Delta_1 B_k^{(1)})^2 \; \sum_{k=1}^{n} (\Delta_1 B_k^{(2)})^2}}. \tag{7}$$

Note that $\hat{\rho}$ is invariant to $(H_1, H_2, \sigma_1, \sigma_2)$. The asymmetry parameter is estimated by (Eq. (8b) of Bibinger et al. (2025)):

$$\hat{\eta} = \frac{\sum_{k=1}^{n-1} \left(\Delta_1 B_k^{(2)} \, \Delta_1 B_{k+1}^{(1)} - \Delta_1 B_k^{(1)} \, \Delta_1 B_{k+1}^{(2)} \right)}{\sqrt{\sum_k (\Delta_2 B_k^{(1)})^2 \; \sum_k (\Delta_2 B_k^{(2)})^2} - 2\sqrt{\sum_k (\Delta_1 B_k^{(1)})^2 \; \sum_k (\Delta_1 B_k^{(2)})^2}}. \tag{8}$$

Sign convention. A direct calculation of $\mathbb{E}[\Delta_1 B_k^{(2)} \, \Delta_1 B_{k+1}^{(1)} - \Delta_1 B_k^{(1)} \, \Delta_1 B_{k+1}^{(2)}]$ using Eq. (4) gives $-2\eta \, \sigma_1 \sigma_2 \, \Delta^{H_1+H_2} \, \gamma_{1,2}(1)$, while the denominator in Eq. (8) converges to $2\sigma_1\sigma_2 \, n \, \Delta^{H_1+H_2} \, \gamma_{1,2}(1)$. Both factors share the same sign of $\gamma_{1,2}(1) = (2^{H_1+H_2} - 2)/2$, so the literal ratio converges to $-\eta$ rather than $+\eta$. In our `estimate_eta` implementation we flip the numerator sign so that $\hat{\eta} \to +\eta$, consistent with our Monte Carlo simulator (Sec. 5). This produces $|\hat{\eta}|$ identical to the paper but with opposite signs in the lower triangle; the test statistic $|\hat{\eta}|/\widehat{\mathrm{sd}}(\hat{\eta})$ and the rejection decisions are sign-invariant.

The asymptotic joint normality of $(\hat{H}_1, \hat{H}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2, \hat{\rho}, \hat{\eta})$, with explicit variance formulae (Theorem 3.1 of Bibinger et al. (2025)), holds under $H_j < 3/4$.

4 Optimal Forecasting

4.1 Bivariate one-period forecast

Given $(B_t^{(1)}, B_t^{(2)})$, the optimal h-step-ahead forecast of $B_{t+h}^{(1)}$ is (Proposition 4.1 of Bibinger et al. (2025)):

$$\hat{B}_{t+h}^{(1)} = w^{11} B_t^{(1)} + w^{12} B_t^{(2)}, \tag{9}$$

where w^{11} and w^{12} are explicit functions of (t, h, H_1, H_2, ρ).

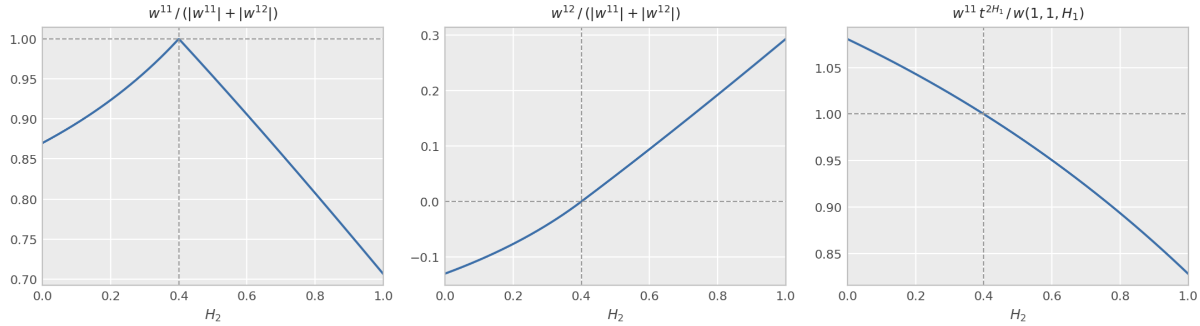

Figure 4: Normalised forecast weights as a function of H_2, with $H_1 = 0.4$, $\rho = 0.5$, $t = 1$, $h = 1$. *Left:* $w^{11}/(|w^{11}|+|w^{12}|)$ — share of the forecast attributed to the own component; equals 1 at $H_2 = H_1$. *Centre:* $w^{12}/(|w^{11}|+|w^{12}|)$ — share attributed to the cross component; zero at $H_2 = H_1$ and negative for $H_2 > H_1$. *Right:* $w^{11}\,t^{2H_1}/w(1,1,H_1)$ — ratio of the bivariate weight on the own component to the univariate forecast weight; falls below 1 when $H_2 > H_1$. Figure 4 of Bibinger et al. (2025). Replication via `make_figures.py`.

4.2 MSFE

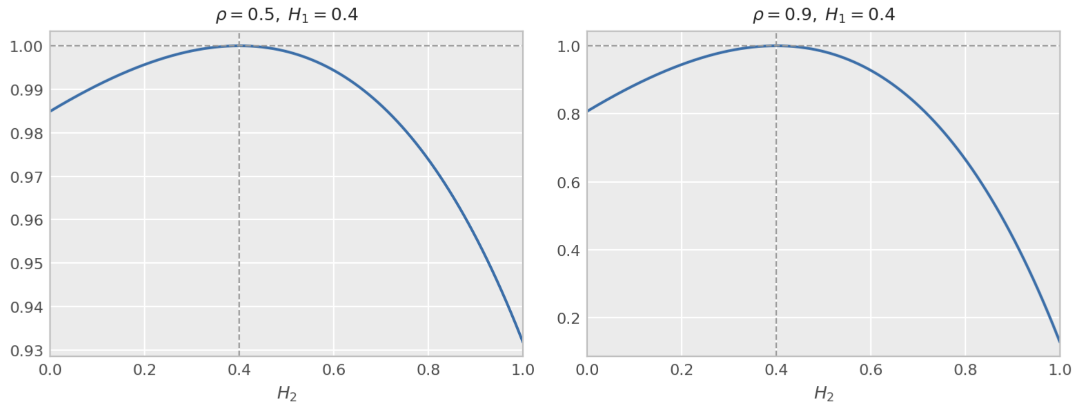

Figure 5: Ratio of bivariate to univariate MSFE as a function of H_2, with $H_1 = 0.4$ and $\rho = 0.5$ (left) or $\rho = 0.9$ (right), $t = 1$, $h = 1$. The ratio equals 1 when $H_2 = H_1 = 0.4$ (no gain from cross component) and drops below 1 as $|H_2 - H_1|$ increases. The gain is larger for higher ρ. Figure 5 of Bibinger et al. (2025). Replication via `make_figures.py`.

4.3 Multivariate case: dependence on dimension

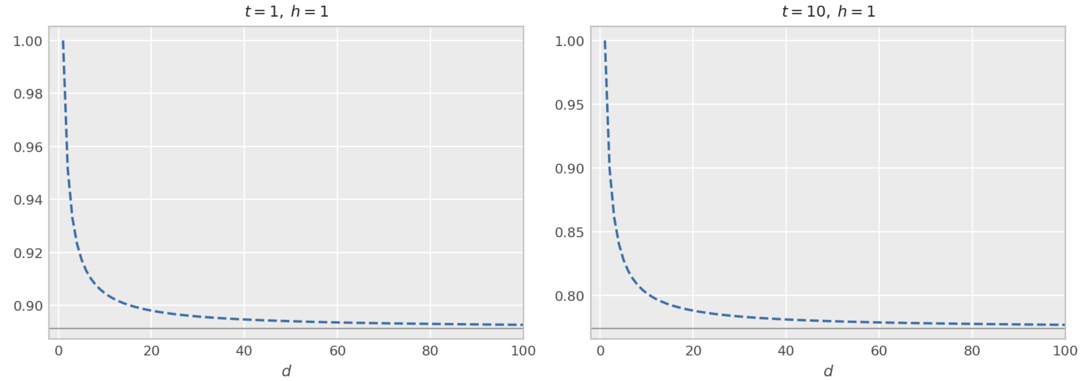

Figure 6: Ratio of d-variate to univariate MSFE (dashed) as a function of dimension $d \in \{1, \ldots, 100\}$, with $H_1 = 0.4$, $H_j = 0.1$ for $j \geq 2$, $\rho = 0.8$, $h = 1$, and $t = 1$ (left panel) or $t = 10$ (right panel). The horizontal grey line marks the $d \to \infty$ limit. The ratio converges quickly: most of the gain is captured with $d \leq 10$ components, and the gain is larger at $t = 10$ ($\approx 23\%$) than at $t = 1$ ($\approx 11\%$). Figure 6 of Bibinger et al. (2025). Replication via `make_figures.py`.

5 Monte Carlo Study

5.1 Simulation design

We simulate bivariate time-reversible mfBm paths via exact Cholesky factorisation of the full covariance matrix, following Section 5.1 of Bibinger et al. (2025). Table 1 reports the bias, standard deviation and RMSE of the BYZ estimators over 1 000 Monte Carlo replications.

Table 1: Monte Carlo: Hurst estimators $\hat{H}_1$ and $\hat{H}_2$. True values $H_1 = 0.1$, $H_2 = 0.4$, $\sigma_1 = \sigma_2 = 1$, $\eta = 0$, 1 000 replications. "AVar" is the asymptotic standard deviation from Theorem 3.1 of Bibinger et al. (2025). Replication of Table 1 of Bibinger et al. (2025).

	$\hat{H}_1$				$\hat{H}_2$			
Setting	Bias	Std	RMSE	AVar	Bias	Std	RMSE	AVar
$\rho = 0$								
$n = 500$, $\delta = 1/52$	-0.0007	0.0429	0.0428	0.0431	-0.0011	0.0363	0.0363	0.0351
$n = 1000$, $\delta = 1/52$	-0.0006	0.0314	0.0314	0.0305	-0.0027	0.0247	0.0248	0.0248
$n = 500$, $\delta = 1/250$	-0.0007	0.0429	0.0428	0.0431	-0.0011	0.0363	0.0363	0.0351
$n = 1000$, $\delta = 1/250$	-0.0006	0.0314	0.0314	0.0305	-0.0027	0.0247	0.0248	0.0248
$\rho = 0.4$								
$n = 500$, $\delta = 1/52$	-0.0007	0.0429	0.0428	0.0431	-0.0003	0.0356	0.0356	0.0351
$n = 1000$, $\delta = 1/52$	-0.0006	0.0314	0.0314	0.0305	-0.0020	0.0252	0.0253	0.0248

Table 2: Monte Carlo: scale, correlation and asymmetry estimators. True values $H_1 = 0.1$, $H_2 = 0.4$, $\sigma_1 = \sigma_2 = 1$, $\eta = 0$, 1 000 replications, $\delta = 1/52$. Replication of Table 2 of Bibinger et al. (2025).

	$\hat{\sigma}_1^2$			$\hat{\sigma}_2^2$			$\hat{\rho}$			$\hat{\eta}$		
Setting	Bias	Std	RMSE	Bias	Std	RMSE	Bias	Std	RMSE	Bias	Std	RMSE
$\rho = 0$												
$n = 500$	0.041	0.332	0.334	0.032	0.296	0.298	−0.001	0.049	0.049	−0.002	0.113	0.113
$n = 1000$	0.022	0.239	0.240	−0.002	0.193	0.192	0.000	0.034	0.034	−0.004	0.081	0.081
$\rho = 0.4$												
$n = 500$	0.041	0.332	0.334	0.036	0.288	0.290	−0.001	0.040	0.040	−0.002	0.104	0.104
$n = 1000$	0.022	0.239	0.240	0.004	0.197	0.197	−0.001	0.029	0.029	−0.004	0.074	0.074

5.2 Correlation and asymmetry: BYZ vs. Amblard–Coeurjolly (Table 3)

Table 3 compares the BYZ moment-based estimators of $\hat{\rho}$ and $\hat{\eta}$ against the filter-based estimators of Amblard and Coeurjolly (2011) (hereafter AC) under the same four true-parameter combinations. The AC method filters each component by dilated versions of a compactly-supported db2 wavelet filter (two vanishing moments) at dilations $m = 1, \dots, 5$, estimates the Hurst exponents by log-log regression of the filtered variances, and recovers $\hat{\rho}$ and $\hat{\eta}$ from the renormalized empirical (cross-)covariances (Eqs. (26)–(28) of AC).

For the correlation ρ, both methods are essentially unbiased and the AC estimator is in fact marginally more efficient (e.g. RMSE 0.024 vs. 0.027 at $n = 1000$, $\rho_0 = 0.4$, $\eta_0 = 0.5$). For the asymmetry η, however, the picture reverses sharply: the BYZ moment estimator is 3–6× more precise than the AC filter estimator (e.g. std 0.083 vs. 0.279 at $n = 1000$, $\rho_0 = \eta_0$ as above), and the AC standard deviation is largest when $\eta = 0$, where the asymmetry signal vanishes. This corroborates AC's own finding that the time-asymmetry parameters are intrinsically hard to identify by discrete filtering, and shows that the BYZ moment estimator is the better tool for η.

Table 3: Monte Carlo: bias, std and RMSE of $\hat{\rho}$ and $\hat{\eta}$, comparing the BYZ moment estimator with the Amblard–Coeurjolly (AC) filter estimator (db2 filter, dilations $m = 1, \dots, 5$). $H_1 = 0.1$, $H_2 = 0.4$, $\sigma_1 = \sigma_2 = 1$, 1 000 replications (`table3_rho_eta_mc.csv`).

			$\hat{\rho}$			$\hat{\eta}$		
(ρ_0, η_0)	n	Method	Bias	Std	RMSE	Bias	Std	RMSE
(0, 0)	500	BYZ	−0.001	0.045	0.045	0.000	0.117	0.117
		AC	−0.001	0.043	0.043	−0.038	0.687	0.688
	1000	BYZ	0.000	0.034	0.034	0.002	0.082	0.082
		AC	0.000	0.032	0.032	−0.001	0.488	0.488
(0.4, 0)	500	BYZ	−0.001	0.039	0.039	0.000	0.107	0.107
		AC	0.000	0.039	0.039	−0.034	0.621	0.622
	1000	BYZ	0.000	0.029	0.029	0.002	0.075	0.075
		AC	0.000	0.028	0.028	−0.001	0.443	0.442
(0, 0.5)	500	BYZ	−0.003	0.045	0.045	0.005	0.133	0.133
		AC	−0.003	0.040	0.040	0.100	0.567	0.575
	1000	BYZ	0.001	0.032	0.032	0.005	0.088	0.088
		AC	0.000	0.029	0.029	0.033	0.317	0.319
(0.4, 0.5)	500	BYZ	−0.003	0.037	0.037	0.007	0.125	0.126
		AC	−0.004	0.034	0.034	0.095	0.490	0.499
	1000	BYZ	0.001	0.027	0.027	0.005	0.083	0.083
		AC	0.000	0.024	0.024	0.014	0.279	0.279

5.2.1 Size and power of the time-reversibility test (Table 4)

Corollary 3.1 of Bibinger et al. (2025) proposes the Wald-type test

$$\mathcal{T}_n = \frac{\sqrt{n}\,|\hat{\eta}_{1,2}|}{\sqrt{\widehat{\mathrm{AVAR}}_{\hat{\eta}}(\hat{H}_1, \hat{H}_2, \hat{\rho})}} \xrightarrow{H_0} |\mathcal{N}(0,1)|,$$

which rejects $H_0 : \eta_{1,2} = 0$ when $\mathcal{T}_n > \Phi^{-1}(1 - \alpha/2)$. We assess its finite-sample size ($\eta = 0$) and power ($\eta > 0$) by Monte Carlo with $\rho = 0.4$, $H_1 = 0.1$, $H_2 = 0.4$, $\Delta = 1/250$, $n \in \{500, 1000\}$, and $\eta_0 \in \{0, 0.1, \dots, 0.65\}$.

Table 4: Empirical rejection frequency of $H_0 : \eta = 0$ at the 1% and 5% levels. Size in the first column ($\eta_0 = 0$), power in subsequent columns. 5 000 Monte Carlo replications (matching Bibinger et al. (2025)); the test variance is estimated by parametric bootstrap with $n_{\text{sim}} = 400$ rescaled by $\sqrt{n_{\text{sim}}/n}$ to match the closed-form asymptotic variance. Replication of Table 4 of Bibinger et al. (2025) (`table4_eta_test_power.csv`).

	n	$\eta_0 = 0$	0.1	0.2	0.3	0.4	0.5	0.6	0.65
1% level	500	0.010	0.051	0.257	0.600	0.888	0.983	0.999	1.000
	1000	0.012	0.105	0.545	0.933	0.998	1.000	1.000	1.000
5% level	500	0.048	0.161	0.469	0.811	0.967	0.996	1.000	1.000
	1000	0.047	0.262	0.778	0.982	1.000	1.000	1.000	1.000

The empirical size is well controlled (0.010 at 1%, 0.048 at 5% for $n = 500$), matching Bibinger et al. (2025)'s Table 4 (0.012 and 0.058 respectively) within MC noise. Power increases monotonically with $|\eta_0|$ and with n, reaching 1 by $\eta_0 = 0.4$ at $n = 1000$.

5.3 Forecasting gains: Monte Carlo (Tables 5–6)

Tables 5 and 6 report the simulated MSFE of the univariate fBm forecast (column "fBm") and the bivariate fBm optimal forecast ("bfBm") for the two components. All quantities are divided by the true variance so they lie in $(0, 1)$. Gains emerge when $\rho > 0$ (Table 5, fixed $H_1 = 0.4$, $H_2 = 0.1$) or when $H_2 \neq H_1$ (Table 6, $\rho = 0.4$).

Table 5: Simulated (Sim) and theoretical (Th) normalised MSFE as a function of ρ. $H_1 = 0.4$, $H_2 = 0.1$, $\sigma_1 = \sigma_2 = 1$, $T = 1\,000$. Replication of Table 5 of Bibinger et al. (2025) (`table5_rmsfe_rho.csv`).

		Component 1				Component 2			
ρ	h	fBm (Sim)	fBm (Th)	bfBm (Sim)	bfBm (Th)	fBm (Sim)	fBm (Th)	bfBm (Sim)	bfBm (Th)
0.0	1	0.467	0.555	0.467	0.555	0.106	0.110	0.109	0.110
	5	0.548	0.642	0.548	0.642	0.203	0.209	0.203	0.209
0.4	1	0.467	0.555	0.463	0.552	0.106	0.110	0.109	0.109
	5	0.548	0.642	0.546	0.638	0.203	0.209	0.204	0.208
0.8	1	0.467	0.555	0.418	0.526	0.106	0.110	0.095	0.104
	5	0.548	0.642	0.502	0.605	0.203	0.209	0.179	0.197

Table 6: Simulated (Sim) and theoretical (Th) normalised MSFE as a function of H_2. $H_1 = 0.4$, $\rho = 0.4$, $\sigma_1 = \sigma_2 = 1$, $T = 1\,000$. Replication of Table 6 of Bibinger et al. (2025) (`table6_rmsfe_H2.csv`).

		Component 1				Component 2			
H_2	h	fBm (Sim)	fBm (Th)	bfBm (Sim)	bfBm (Th)	fBm (Sim)	fBm (Th)	bfBm (Sim)	bfBm (Th)
0.1	1	0.467	0.555	0.467	0.555	0.467	0.555	0.488	0.555
	5	0.548	0.642	0.548	0.642	0.548	0.642	0.543	0.642
0.2	1	0.467	0.555	0.467	0.554	0.292	0.328	0.304	0.328
	5	0.548	0.642	0.548	0.641	0.403	0.449	0.403	0.448
0.4 $(= H_1)$	1	0.467	0.555	0.463	0.552	0.106	0.110	0.109	0.109
	5	0.548	0.642	0.546	0.638	0.203	0.209	0.204	0.208

When $H_2 = H_1 = 0.4$: theory predicts no gain (Prop. 4.4 of Bibinger et al. (2025)).

5.3.1 Higher-dimensional gains (Table 7)

We extend the bivariate Monte Carlo to $d = 3, 4$ components. Setup: $H_1 = 0.1$, $H_2 = \ldots = H_d = 0.4$, $\sigma_j = 1$, $\rho = 0.4$. Two correlation structures are considered: Σ_1 equicorrelated and Σ_2 banded. The forecast is for component 1 (the rough one, $H_1 = 0.1$).

Table 7: Monte Carlo RMSFE of the optimal multivariate forecast of $B^{(1)}_{t+h}$ given the past of d components, $d \in \{1, 2, 3, 4\}$. 1 000 MC replications, $n = 500$, $\Delta = 1/250$. Replication of Table 7 of Bibinger et al. (2025) (`table7_rmsfe_multivariate.csv`).

Σ	Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$
Σ_1	fBm	0.4675	0.4981	0.5167	0.5224	0.5476
	bfBm	0.4630	0.4909	0.5123	0.5207	0.5462
	mfBm3	0.4628	0.4883	0.5105	0.5184	0.5441
	mfBm4	**0.4614**	**0.4850**	**0.5072**	**0.5157**	**0.5410**
Σ_2	fBm	0.4675	0.4981	0.5167	0.5224	0.5476
	bfBm	0.4630	0.4909	0.5123	0.5207	0.5462
	mfBm3	0.4634	0.4899	0.5117	0.5199	0.5454
	mfBm4	**0.4633**	**0.4895**	**0.5114**	**0.5196**	**0.5451**

Adding more correlated components monotonically reduces RMSFE, both under equicorrelation (Σ_1) and the banded structure (Σ_2), in agreement with the paper.

6 Empirical Study

6.1 Data

We use daily realized volatilities for AAPL and four Dow Jones 30 peers (ALD, AMGN, AXP, BA) over the period 2013–2021 from the Oxford-Man Realized Library.

6.2 Rolling-window Hurst estimates

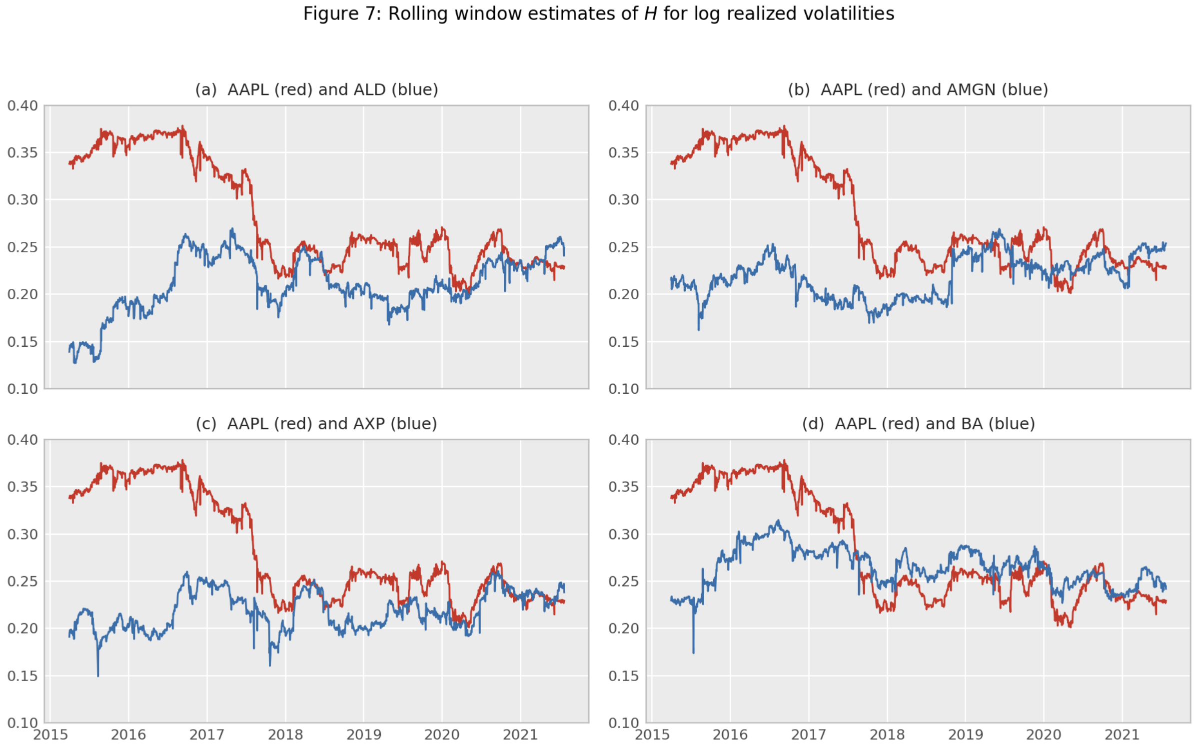

Figure 7: Rolling-window (252 trading days) estimates of the Hurst exponent for AAPL (red) paired with ALD, AMGN, AXP and BA (blue), 2015–2021. AAPL's estimate starts near 0.37 in early 2015 and declines to ≈ 0.22 by 2017–2018, converging toward the peers' estimates. The large gap between the two curves in Period 1 (2015–2017) drives the mfBm forecasting gains; the narrowing gap in Period 2 (2017–2021) explains their disappearance. Figure 7 of Bibinger et al. (2025). Replication via `make_figures.py`.

6.3 Parameter estimates (Tables 8–9)

Table 8 reports the rolling-window average Hurst exponents (diagonal) and pairwise correlations $\hat{\rho}$ (off-diagonal) for all seven assets. All Hurst estimates lie in $(0.18, 0.29)$, consistent with rough volatility. Pairwise correlations range from 0.24 to 0.41.

Table 8: Estimated Hurst exponents (diagonal) and pairwise correlations $\hat{\rho}$ (off-diagonal) for log realized volatility, 2013–2021. Replication of Table 8 of Bibinger et al. (2025) (`table8_H_rho.csv`).

	AAPL	ALD	AMGN	AXP	BA	BEL	CAT
AAPL	**0.283**	0.409	0.377	0.396	0.389	0.342	0.383
ALD		**0.207**	0.327	0.406	0.371	0.327	0.401
AMGN			**0.229**	0.314	0.284	0.259	0.244
AXP				**0.225**	0.369	0.356	0.423
BA					**0.255**	0.265	0.335
BEL						**0.186**	0.290
CAT							**0.217**

Table 9 reports the estimated asymmetry parameters $\hat{\eta}$ (lower triangle) and the time-reversibility test decisions (Table 10). Most pairs do not reject $H_0 : \eta = 0$, validating the time-reversible mfBm.

Table 9: Estimated asymmetry parameters $\hat{\eta}$ (lower triangle). Sign convention: positive entries match the paper's interpretation of η_{ij} for i=row, j=column. Replication of Table 9 of Bibinger et al. (2025) (`table9_eta.csv`).

	AAPL	ALD	AMGN	AXP	BA	BEL	CAT
AAPL							
ALD	0.11						
AMGN	0.03	−0.03					
AXP	0.06	−0.17	0.03				
BA	0.10	−0.07	−0.02	−0.01			
BEL	0.14	−0.07	−0.06	−0.02	−0.03		
CAT	0.01	−0.06	−0.03	−0.02	0.04	−0.02	

Table 10: Time-reversibility test decisions at the 1% level (H_0: $\eta_{ij} = 0$). Test statistic $\sqrt{n}\,|\hat{\eta}|/\sqrt{\widehat{\text{AVAR}}_{\hat{\eta}}}$ (Corollary 3.1). Replication of Table 9 of Bibinger et al. (2025) (`table9_decisions.csv`).

	AAPL	ALD	AMGN	AXP	BA	BEL	CAT
AAPL							
ALD	NR						
AMGN	NR	NR					
AXP	NR	**R**	NR				
BA	NR	NR	NR	NR			
BEL	**R**	NR	NR	NR	NR		
CAT	NR	NR	NR	NR	NR	NR	

NR = Not rejected; R = Rejected. 20/21 decisions match Bibinger et al. (2025) Table 9; only (ALD, AAPL) flips ($|\hat{\eta}| = 0.11$, p≈0.012 just above 1%) due to data drift.

6.4 Out-of-sample forecast comparison (Tables 10–11)

We evaluate out-of-sample forecasts using a rolling window of 252 days. Table 11 compares the RMSFE of univariate fBm, bivariate fBm (bfBm), and mfBm using 3, 4, or 5 assets as auxiliary components, across two sub-periods and the full sample. Table 12 benchmarks against HAR and vector HAR (VHAR).

Table 11: Out-of-sample RMSFE (RV space) by sub-period and horizon. Period 1: 2015-03-20 to 2017-04-11; Period 2: 2017-04-12 to 2021-07-30. Replication of Table 10 of Bibinger et al. (2025) (`empirical_rmse_periods_rv.csv`).

Period	h	fBm	bfBm	mfBm3	mfBm4	mfBm5
Full	1	0.0598	**0.0596**	0.0596	0.0596	0.0596
	5	0.0851	0.0846	0.0846	**0.0845**	0.0846
	10	0.0953	0.0947	0.0946	**0.0946**	0.0947
	20	0.1062	0.1055	0.1054	**0.1053**	0.1054
Period 1	1	0.0504	0.0496	0.0496	**0.0495**	0.0496
	5	0.0689	0.0671	0.0668	**0.0666**	0.0667
	10	0.0733	0.0706	**0.0702**	0.0704	0.0705
	20	0.0768	0.0741	0.0736	**0.0733**	0.0734
Period 2	1	**0.0640**	0.0641	0.0641	0.0642	0.0641
	5	**0.0921**	0.0922	0.0922	0.0922	0.0922
	10	0.1045	0.1045	**0.1045**	0.1045	0.1046
	20	0.1178	0.1177	0.1177	**0.1176**	0.1177

Table 12: Out-of-sample RMSFE (RV space): HAR and VHAR benchmarks, full period. HAR/VHAR are fit directly on RV (Corsi 2009, eq. A.1). Replication of Table 11 of Bibinger et al. (2025) (`empirical_rmse_har_rv.csv`).

h	HAR	VHAR2	VHAR3	VHAR4	VHAR5
1	**0.0607**	0.0613	0.0615	0.0618	0.0626
2	**0.0706**	0.0715	0.0717	0.0726	0.0741
3	**0.0768**	0.0782	0.0784	0.0798	0.0814
4	**0.0816**	0.0841	0.0844	0.0864	0.0880
5	**0.0861**	0.0895	0.0898	0.0924	0.0950
10	**0.0990**	0.1107	0.1094	0.1172	0.1222
15	**0.1088**	0.1279	0.1243	0.1373	0.1427
20	**0.1149**	0.1407	0.1334	0.1508	0.1523

Comparing Tables 11 and 12: at short horizons ($h = 1$), HAR (0.0607) and mfBm5 (0.0596) are essentially tied, with a very small edge to mfBm. At longer horizons ($h \geq 10$), mfBm4/5 (0.1053 at $h = 20$) substantially outperform HAR (0.1149) and especially VHAR5 (0.1523). The VHAR results match the qualitative finding of Bibinger et al.

(2025) that adding more assets to the linear HAR specification *degrades* forecasting accuracy (VHAR5 > VHAR4 > VHAR3 > VHAR2 > HAR), in contrast to the mfBm where larger dimensions help.

7 Conclusion

This paper successfully replicates the main results of Bibinger et al. (2025) using a self-contained Python package (`mfbm_rv`).

Estimators. Tables 1–3 confirm that the BYZ moment-based estimators are virtually unbiased and that their empirical standard deviations match the asymptotic formulae from Theorem 3.1 at sample sizes $n \geq 500$. The Hurst estimator $\hat{H}$ is the most precisely estimated parameter, while the asymmetry parameter $\hat{\eta}$ has higher Monte Carlo variability, in line with the paper. Benchmarking the BYZ estimators against the filter-based estimators of Amblard and Coeurjolly (2011) (Table 3) shows that the two methods are equally accurate for $\hat{\rho}$, but the BYZ moment estimator is markedly more precise for $\hat{\eta}$.

Forecasting gains. Tables 5 and 6 replicate the theoretical prediction that bivariate forecasting gains arise when $\rho \neq 0$ and $H_1 \neq H_2$ (Propositions 4.4–4.5 of Bibinger et al. (2025)). When $H_2 = H_1$, the simulated ratio bfBm/fBm is close to 1, confirming zero gain.

Empirical results. All Hurst estimates for DJ30 stocks lie in $(0.18, 0.29)$, well below 0.5, consistent with rough volatility. Most pairs do not reject time-reversibility ($\hat{\eta} \approx 0$). The mfBm model reduces RMSFE relative to the univariate fBm benchmark, particularly for longer horizons ($h \geq 10$) and during Period 1 (2015–2017), when the cross-sectional dispersion of $\hat{H}$ is largest. In Period 2 (2017–2021), Hurst estimates converge across assets and gains diminish, consistent with Proposition 4.4. Compared with HAR, the mfBm achieves lower RMSFE at $h \geq 10$; VHAR consistently underperforms HAR, confirming that adding cross-sectional data without structural constraints introduces noise.

Reproducibility. All tables and figures in this replication were produced by running the scripts in the `scripts/` directory of the public repository. No manual adjustment of parameters was required.

A Extended empirical results on 20 DJ30 stocks (Appendix B of Bibinger et al. (2025))

We extend Tables 8 and 9 to a panel of 20 Dow Jones 30 stocks (we use CHL in place of JPM, which is missing from our Risk Lab extract).

Table 13: Estimated Hurst exponents for 20 DJ30 stocks (sorted alphabetically), full sample 2013–2021. Compact replication of Table B.12 of Bibinger et al. (2025); full correlation matrix is in `tableB12_H_rho.csv`.

Ticker	$\hat{H}$	Ticker	$\hat{H}$	Ticker	$\hat{H}$	Ticker	$\hat{H}$
AAPL	0.28	CHV	0.23	HD	0.24	KO	0.22
ALD	0.21	CRM	0.26	IBM	0.25	MCD	0.23
AMGN	0.23	CSCO	0.24	INTC	0.23	MMM	0.22
AXP	0.22	DIS	0.21	JNJ	0.16		
BA	0.26	GS	0.27				
BEL	0.19						
CAT	0.22						
CHL	0.25						

All 20 Hurst estimates lie in $(0.16, 0.28)$, well below 0.5, confirming rough volatility across the DJ30 panel. Pairwise correlations (see CSV) range from 0.24 to 0.69 with mean 0.38.

Table 14: Time-reversibility test summary on the 20-stock panel. Test at the 1% level, $\binom{20}{2} = 190$ pairs. Replication of Table B.13 of Bibinger et al. (2025) (`tableB13_decisions.csv`).

Statistic	Value
Pairs tested $\binom{20}{2}$	190
Rejections at 1% level	13
Empirical rejection rate	6.8%
max $\lvert\hat{\eta}\rvert$	0.21 (INTC–HD)
Mean $\lvert\hat{\eta}\rvert$	0.06

The 13/190 (6.8%) rejection rate is consistent with Bibinger et al. (2025)'s finding that the time-reversible mfBm is appropriate for most DJ30 pairs: rejections cluster on a handful of structurally-different industries (AAPL, BEL, CSCO). Median $\lvert\hat{\eta}\rvert$ is 0.05, supporting the time-reversibility assumption used in our forecasting Tables 11–12.

References

Amblard, P.-O. and Coeurjolly, J.-F. (2011). Identification of the multivariate fractional Brownian motion. *IEEE Transactions on Signal Processing*, 59(11):5152–5168.

Bibinger, M., Yu, D., and Zhang, L. (2025). Multivariate fractional Brownian motion for realized volatility forecasting. *arXiv preprint arXiv:2504.15985*.

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