LATENT VOLATILITY CONTAGION IN ROUGH VOLATILITY MODELS
SPECTRAL THRESHOLDS AND DETECTABILITY

J. VIDAL LLAURADÓ

Abstract. Short-maturity option prices and physical path laws need not observe the same directed rough-volatility channel. The paper separates two projections of a latent cross-asset volatility component. The pricing projection records whether the cross-kernel enters the leading Bachelier skew coefficient. The path-law projection records the residual Volterra perturbation after covariance normalisation. In a bivariate rough Volterra model, asset Y is driven by an idiosyncratic kernel with Hurst exponent H_Y and by a directed contagion kernel from asset X with exponent H_{XY}.

Two thresholds govern the separation. The *pricing threshold* $H_{XY} = H_Y$ determines whether the contagion component contributes to the leading short-maturity Bachelier skew. In the smoothing regime $H_{XY} > H_Y$, the *path-space threshold* $H_{XY} = H_Y + \frac{1}{4}$ determines whether the coupled and uncoupled Gaussian volatility laws are mutually singular or equivalent. The interval $H_Y < H_{XY} \leq H_Y + \frac{1}{4}$ is the *latent contagion regime*. There the directed channel is absent from the leading pricing projection but remains detectable in the physical path law.

Smooth non-degenerate modulation leaves this threshold geometry unchanged. The relative covariance operator reduces to a principal fractional factor, and its singular-value order gives the Feldman–Hájek dichotomy. The result identifies a rough-volatility regime in which pricing visibility stops before path-law detectability.

1. Introduction

Rough-volatility models give a tractable language for fractional short-maturity scaling and Volterra memory under both the physical and pricing measures [1, 2, 3, 4, 5]. They do not, by themselves, decide whether a directed cross-asset volatility route is visible to options, to the physical path law, or to both. The observation experiment matters.

The analysis separates two projections of a directed Volterra volatility channel. The short-maturity pricing projection records whether the cross-kernel enters the leading skew coefficient. The path-law projection records whether the coupled Gaussian volatility law is equivalent to, or singular with respect to, the uncoupled law. These projections have different thresholds.

Cross-asset spillover work supplies the empirical pressure for the distinction. Rough microstructural limits generate multivariate rough volatility with nontrivial cross-dependence [6]. Multivariate fractional models with heterogeneous exponents produce asymmetric cross-covariance effects and forecasting gains [7, 8, 9]. Contagion and connectedness studies also warn that stronger comovement need not identify a latent directed channel by itself [10, 11, 12, 13]. The comparison here is narrower and more formal: option-price visibility is separated from Gaussian path-law detectability for one directed rough Volterra component.

In the bivariate model, asset Y has idiosyncratic exponent H_Y and receives a contagion component from asset X with exponent H_{XY}. The pricing threshold is $H_{XY} = H_Y$. Below this boundary, the directed component contributes to the leading short-maturity ATM skew in the model's Bachelier expansion. The path-law threshold is $H_{XY} = H_Y + \frac{1}{4}$ in the smoothing regime $H_{XY} > H_Y$. It separates mutual singularity from equivalence of the coupled and uncoupled Gaussian path laws. Between them lies the regime isolated here,

$$H_Y < H_{XY} \leq H_Y + \tfrac{1}{4},$$

Date: March 2026.

2020 *Mathematics Subject Classification.* 60G15, 60G22, 91G20, 60H07.

Key words and phrases. rough volatility, latent volatility contagion, Gaussian Volterra models, option-price visibility, path-space detectability, Feldman–Hájek dichotomy, Hilbert–Schmidt operators.

DOI: 10.2139/ssrn.6520718.

where contagion is absent from the leading pricing projection and still remains detectable in the physical path law.

The path-space comparison is a Gaussian-law comparison. The Feldman–Hájek dichotomy [14, 15, 16, 17, 18, 19, 20] reduces it to a Hilbert–Schmidt condition on a relative covariance operator. The singular-value order of that operator is determined by the fractional Volterra backbone, using the asymptotics of Birman and Solomyak [21], Dostanić [22], Tuan and Gorenflo [23], and Gorenflo and Tuan [24]. Smooth non-degenerate modulation preserves the same threshold beyond the pure fractional benchmark used in the classical theory.

The threshold statement is static. It identifies the smoothing latent-contagion regime before sequential revelation, observable filtrations, or feasible high-frequency inference are imposed.

The model class uses modulated Volterra kernels $(t-s)^{H-1/2}L(t,s)$ with C^1 modulating functions, including time-inhomogeneous and mean-reverting kernels used in rough-volatility practice. The operator proof separates each kernel into a Riemann–Liouville principal part and a diagonal-vanishing Volterra remainder, then transfers the singular-value asymptotics to the relative covariance operator. A bivariate rough Bergomi-type specification with kernel $(t-s)^{H-1/2}e^{-\lambda(t-s)}$ serves as the benchmark chart.

Main thresholds. The threshold comparison has three components. The pricing threshold $H_{XY} = H_Y$ determines whether contagion enters the leading short-maturity skew. The path-space detectability threshold $H_{XY} = H_Y + \frac{1}{4}$ determines, in the smoothing regime, whether the coupled and uncoupled Gaussian laws are equivalent or mutually singular.

The latent contagion regime lies between those boundaries. In that interval, a cross-asset channel disappears from leading short-maturity option asymptotics and remains detectable under the physical law. This is a smoothing-regime statement: $H_{XY} > H_Y$. Rougher directed channels belong to a separate observable-screening problem.

Modulation is the third issue. For the modulated Volterra kernels used in practice, smooth non-degenerate modulation leaves unchanged the spectral order that drives the threshold comparison.

Organization. Section 2 defines the bivariate rough Volterra model. Sections 3 through 6 give the operator reduction and derive the spectral threshold. Sections 7 and 8 translate that threshold into pricing visibility and path-space detectability. Sections 9 and 10 state the latent contagion regime, Section 11 records finite-dimensional threshold diagnostics, and Sections 12–13 state the dynamic and inference boundaries.

2. Model

Let $(\Omega, \mathcal{F}, \mathbb{P})$ carry independent standard Brownian motions W^X and $W^\perp$. When pricing is discussed, $\mathbb{Q}$ denotes a risk-neutral measure under which the same Gaussian volatility specification is used.

Assumption 2.1 (Kernel regularity and non-degeneracy)**.** Let

$$\Delta_T := \{(t,s) \in [0,T]^2 : 0 \le s \le t \le T\}.$$

For $i \in \{Y, XY\}$ let $L_i : \Delta_T \to \mathbb{R}$ satisfy:

(i) L_i extends to a C^1 function on a neighbourhood of Δ_T;
(ii) L_i and its first partial derivatives are bounded on Δ_T;
(iii) there exists $c_i > 0$ such that

$$L_i(t,t) \ge c_i, \qquad t \in [0,T].$$

Remark 2.2. The diagonal non-degeneracy condition keeps the leading fractional singularity present in both the idiosyncratic and cross kernels. A vanishing diagonal would change the leading singularity and the spectral order governing the threshold. The C^1 condition isolates the diagonal trace. The fractional gain of the diagonal-vanishing remainder is recorded in Proposition 6.1. The pure power-law case is $L_Y \equiv L_{XY} \equiv 1$.

For $i \in \{Y, XY\}$, write

$$K_i(t,s) := (t-s)^{H_i - 1/2} L_i(t,s) \mathbf{1}_{\{0 \le s < t \le T\}}.$$

The idiosyncratic and contagion volatility processes are

$$V_t = \nu_Y \int_0^t (t-s)^{H_Y - 1/2} L_Y(t,s)\, dW_s^{\perp}, \tag{1}$$

$$Z_t = \int_0^t (t-s)^{H_{XY} - 1/2} L_{XY}(t,s)\, dW_s^X, \tag{2}$$

with $H_Y, H_{XY} \in (0, \frac{1}{2})$, $\nu_Y > 0$, and $\eta_{X\to Y} \ge 0$. Set $c = \eta_{X\to Y}/\nu_Y$. The case $c > 0$ is the non-degenerate contagion case; if $c = 0$, the coupled and uncoupled volatility laws coincide and the path-law thresholds below are void. The log-volatility of Y is

$$\log \sigma_t^Y = \log \sigma_0^Y + V_t + c\,\nu_Y\, Z_t - C(t), \tag{3}$$

where $C(t)$ is the convexity correction ensuring $\mathbb{E}_{\mathbb{P}}[\sigma_t^Y] = \sigma_0^Y$. For pricing, the Bachelier dynamics under $\mathbb{Q}$, with $\rho \in (-1, 1)$, are

$$dY_t = \sigma_t^Y (\rho\, dW_t^X + \sqrt{1-\rho^2}\, dW_t^{\perp}). \tag{4}$$

Let μ_η (resp. μ_0) be the law of $\{V_t + c\,\nu_Y Z_t\}_{t\in[0,T]}$ (resp. $\{V_t\}_{t\in[0,T]}$) on $C([0,T])$. Both are centred Gaussian measures. Barred notation $\bar\mu_\eta, \bar\mu_0$ refers to the corresponding Gaussian laws on $L^2([0,T])$; unbarred notation μ_η, μ_0 refers to the induced path-space laws on $C([0,T])$.

Remark 2.3 (Numeraire invariance). The Bachelier specification is adopted for algebraic convenience; the $T^{H-1/2}$ scaling of the ATM skew is invariant to the choice of price numeraire in the short-maturity limit (Fukasawa [25], §2.1).

Covariance operators. Define the modulated Volterra operators $K_Y, K_{XY} : L^2[0,T] \to L^2[0,T]$ by

$$\begin{aligned} (K_Y f)(t) &= \int_0^t (t-s)^{H_Y - 1/2} L_Y(t,s)\, f(s)\, ds, \\ (K_{XY} f)(t) &= \int_0^t (t-s)^{H_{XY} - 1/2} L_{XY}(t,s)\, f(s)\, ds. \end{aligned} \tag{5}$$

The processes V and Z from (1)–(2) are the images of white noise under $\nu_Y K_Y$ and K_{XY} respectively. Their covariance operators on $L^2[0,T]$ are

$$C_Y = \nu_Y^2\, K_Y K_Y^*, \qquad C_{XY} = c^2 \nu_Y^2\, K_{XY} K_{XY}^*. \tag{6}$$

The covariance of the coupled process $V + c\nu_Y Z$ is $C_Y + C_{XY}$. The symbols $K_{0,Y}$ and $K_{0,XY}$ denote the corresponding *pure fractional* (Riemann–Liouville) operators with $L \equiv 1$:

$$(K_{0,i} f)(t) = \int_0^t (t-s)^{H_i - 1/2} f(s)\, ds, \qquad i \in \{Y, XY\}. \tag{7}$$

Remark 2.4 (Hilbert-space setting). The covariance operators act on $L^2[0,T]$. The first comparison is made for centred Gaussian measures on $L^2[0,T]$, where the relative covariance operator is defined. Equivalence of the induced path laws on $C([0,T])$ is then obtained by a transfer proposition using almost-sure continuity of the Volterra processes and the Borel isomorphism induced by the canonical inclusion $C([0,T]) \hookrightarrow L^2([0,T])$.

Continuity and path-space transfer.

Lemma 2.5 (Continuous versions of the modulated Volterra processes)**.** *Let*

$$X_t := \int_0^t (t-s)^{H-\frac{1}{2}} L(t,s)\, dW_s, \qquad t \in [0,T],$$

where $H \in (0, \frac{1}{2})$ and L satisfies Assumption 2.1. Then X admits a continuous modification. In particular, the processes V, Z, and $V + c\nu_Y Z$ admit continuous modifications.

Proof. Fix $0 \le u < t \le T$. By Itô isometry,

$$\mathbb{E}|X_t - X_u|^2 = \int_0^u \big[(t-s)^{H-\frac{1}{2}}L(t,s) - (u-s)^{H-\frac{1}{2}}L(u,s)\big]^2 ds + \int_u^t (t-s)^{2H-1}L(t,s)^2\, ds.$$

The second term is bounded by $C|t-u|^{2H}$ because L is bounded on Δ_T.

For the first term, write

$$\begin{aligned}&(t-s)^{H-\frac{1}{2}}L(t,s) - (u-s)^{H-\frac{1}{2}}L(u,s)\\ &\quad= \big[(t-s)^{H-\frac{1}{2}} - (u-s)^{H-\frac{1}{2}}\big]L(u,s) + (t-s)^{H-\frac{1}{2}}\big[L(t,s) - L(u,s)\big].\end{aligned}$$

Since L and $\partial_t L$ are bounded, the second summand contributes at most

$$C|t-u|^2 \int_0^u (t-s)^{2H-1}\, ds \le C|t-u|^2 \le C_T|t-u|^{2H},$$

because $2H < 1$ and $|t-u| \le T$. For the first summand, the pure fractional estimate for the Riemann–Liouville kernel yields

$$\int_0^u \big[(t-s)^{H-\frac{1}{2}} - (u-s)^{H-\frac{1}{2}}\big]^2 ds \le C_H|t-u|^{2H};$$

see, for instance, the increment estimates for Riemann–Liouville fractional processes in [26, Ch. 2]. Hence

$$\mathbb{E}|X_t - X_u|^2 \le C|t-u|^{2H}.$$

Because $X_t - X_u$ is Gaussian, all moments are controlled by the variance, so in particular

$$\mathbb{E}|X_t - X_u|^4 \le C|t-u|^{4H}.$$

Choose any integer $m > 1/H$. Then Gaussian moment estimates give

$$\mathbb{E}|X_t - X_u|^{2m} \le C_m|t-u|^{2mH},$$

and $2mH > 1$. Kolmogorov's continuity theorem therefore yields a continuous modification of X.

Applying this with $(H, L) = (H_Y, L_Y)$ and $(H, L) = (H_{XY}, L_{XY})$ gives continuous versions of V and Z, hence also of $V + c\nu_Y Z$. □

Proposition 2.6 (Transfer from L^2 laws to path laws). *Let $\bar{\mu}_\eta$ and $\bar{\mu}_0$ denote the laws on $L^2[0,T]$ of the continuous modifications of $V + c\nu_Y Z$ and V, viewed as L^2-valued random variables. Let μ_η and μ_0 be their laws on $C([0,T])$, and let*

$$J : C([0,T]) \to L^2[0,T], \qquad Jx = x,$$

be the canonical inclusion. Then:

(i) J is continuous and injective, and $J(C([0,T]))$ is a Borel subset of $L^2[0,T]$.
(ii) The inverse map $J^{-1} : J(C([0,T])) \to C([0,T])$ is Borel measurable.
(iii) $\bar{\mu}_\eta$ and $\bar{\mu}_0$ are supported on $J(C([0,T]))$, and

$$\mu_\eta = (J^{-1})_\# \bar{\mu}_\eta, \qquad \mu_0 = (J^{-1})_\# \bar{\mu}_0.$$

(iv) The measures μ_η and μ_0 on $C([0,T])$ are equivalent if and only if $\bar{\mu}_\eta$ and $\bar{\mu}_0$ are equivalent on $L^2[0,T]$, and they are mutually singular if and only if $\bar{\mu}_\eta$ and $\bar{\mu}_0$ are mutually singular.

Proof. The continuity of J follows from $\|x\|_{L^2} \le T^{1/2}\|x\|_\infty$, and injectivity is obvious. Since $C([0,T])$ and $L^2([0,T])$ are Polish spaces, the Lusin–Souslin theorem implies that the image $J(C([0,T]))$ is Borel in $L^2([0,T])$ and that J is a Borel isomorphism from $C([0,T])$ onto its image; this is a standard consequence of the Lusin–Souslin theorem. This proves (i) and (ii).

By Lemma 2.5, the processes admit continuous modifications, so their L^2-valued laws assign mass one to $J(C([0,T]))$. For any Borel set $A \subseteq C([0,T])$,

$$\bar{\mu}_\eta(J(A)) = \mathbb{P}(J(V + c\nu_Y Z) \in J(A)) = \mathbb{P}(V + c\nu_Y Z \in A) = \mu_\eta(A),$$

and similarly for μ_0. Thus

$$\bar{\mu}_\eta = J_\# \mu_\eta, \qquad \bar{\mu}_0 = J_\# \mu_0, \qquad \mu_\eta = (J^{-1})_\# \bar{\mu}_\eta, \qquad \mu_0 = (J^{-1})_\# \bar{\mu}_0$$

on the common full-measure support $J(C([0,T]))$.

The null sets are therefore transported exactly by the Borel isomorphism J. If $\bar{\mu}_\eta \ll \bar{\mu}_0$ and $\mu_0(A)=0$, then $\bar{\mu}_0(J(A))=0$, hence $\bar{\mu}_\eta(J(A))=0$ and $\mu_\eta(A)=0$. Conversely, if $\mu_\eta \ll \mu_0$ and $B \subset L^2[0,T]$ is Borel with $\bar{\mu}_0(B)=0$, then $A := J^{-1}(B \cap J(C([0,T])))$ satisfies $\mu_0(A)=0$, so $\bar{\mu}_\eta(B)=\bar{\mu}_\eta(B\cap J(C([0,T])))=\mu_\eta(A)=0$. This proves equivalence in both directions. The same correspondence sends a separating full-measure set in one space to a separating full-measure set in the other, so mutual singularity is preserved as well. This proves (iv). $\square$

3. Operator-theoretic structure of modulated Volterra kernels

3.1. Decomposition of modulated operators.

Lemma 3.1 (Decomposition of a modulated Volterra operator)**.** *Let $H \in (0,\frac{1}{2})$ and define*

$$(K_L f)(t) := \int_0^t (t-s)^{H-\frac{1}{2}} L(t,s) f(s)\, ds, \qquad f \in L^2[0,T],$$

where L satisfies Assumption 2.1. Then

$$K_L = M_L K_0 + V,$$

where

$$(K_0 f)(t) := \int_0^t (t-s)^{H-\frac{1}{2}} f(s)\, ds, \qquad (M_L f)(t) := L(t,t) f(t),$$

and V is a Volterra operator with kernel

$$k_V(t,s) = (t-s)^{H+\frac{1}{2}} R(t,s), \qquad R \in L^\infty(\Delta_T).$$

In particular, M_L is bounded and boundedly invertible on $L^2[0,T]$.

Proof. Write

$$L(t,s) = L(t,t) + \big(L(t,s) - L(t,t)\big).$$

The first term gives $M_L K_0$. For the remainder, since L extends to a C^1 function,

$$L(t,s) - L(t,t) = -\int_s^t \partial_2 L(t,u)\, du.$$

Hence

$$(t-s)^{H-\frac{1}{2}}\big(L(t,s) - L(t,t)\big) = (t-s)^{H+\frac{1}{2}} R(t,s),$$

where

$$R(t,s) := -(t-s)^{-1} \int_s^t \partial_2 L(t,u)\, du.$$

The boundedness of $\partial_2 L$ implies $R \in L^\infty(\Delta_T)$. The bounded invertibility of M_L follows from the diagonal lower bound in Assumption 2.1. $\square$

3.2. Singular value dominance.

Lemma 3.2 (Fractional-order dominance)**.** *Let $K_L = M_L K_0 + V$ be as in Lemma 3.1. Then*

$$s_n(K_L) = O\big(n^{-(H+\frac{1}{2})}\big), \qquad s_n(V) = O\big(n^{-(H+\frac{3}{2})}\big).$$

Consequently, V is of strictly smaller singular-value order than the leading term $M_L K_0$.

Proof. By Birman–Solomyak [21],

$$s_n(K_0) \asymp n^{-(H+\frac{1}{2})}.$$

Since M_L is bounded,

$$s_n(M_L K_0) \le \|M_L\|\, s_n(K_0) = O\big(n^{-(H+\frac{1}{2})}\big).$$

The remainder V has kernel of the form $(t-s)^{H+1/2} R(t,s)$ with $R \in L^\infty(\Delta_T)$, so another application of Birman–Solomyak yields

$$s_n(V) = O\big(n^{-(H+\frac{3}{2})}\big).$$

Applying Ky Fan's inequality with $m = n = \lfloor N/2 \rfloor + 1$,

$$s_N(K_L) \le s_m(M_L K_0) + s_n(V) = O\big(N^{-(H+\frac{1}{2})}\big) + O\big(N^{-(H+\frac{3}{2})}\big),$$

which gives the claimed upper bound for $s_N(K_L)$. □

4. Fractional normalisation

Proposition 4.1 (Fractional normalisation of covariance)**.** *Let*

$$C_Y = \nu_Y^2 K_Y K_Y^*, \qquad C_{0,Y} := K_{0,Y} K_{0,Y}^*.$$

Then there exist bounded multiplication operators M_1, M_2, with M_1 and M_2 boundedly invertible, and a compact operator E_0 such that

$$C_Y = M_1\, C_{0,Y}\, M_2 + E_0,$$

where E_0 is of strictly smaller singular-value order than $C_{0,Y}$. In particular,

$$\lambda_n(C_Y) \asymp n^{-(2H_Y+1)}.$$

Proof. By Lemma 3.1,

$$K_Y = M_Y K_{0,Y} + V_Y,$$

with M_Y boundedly invertible and

$$s_n(V_Y) = O\big(n^{-(H_Y+\frac{3}{2})}\big).$$

Therefore

$$C_Y = \nu_Y^2 (M_Y K_{0,Y} + V_Y)(M_Y K_{0,Y} + V_Y)^* = \nu_Y^2 M_Y C_{0,Y} M_Y^* + E_0,$$

where

$$E_0 := \nu_Y^2 \big(M_Y K_{0,Y} V_Y^* + V_Y K_{0,Y}^* M_Y^* + V_Y V_Y^*\big).$$

Each term in E_0 contains at least one smoother factor V_Y or V_Y^*. The estimates are

$$s_n(M_Y K_{0,Y} V_Y^*) \le \|M_Y K_{0,Y}\|\, s_n(V_Y) = O\big(n^{-(H_Y+\frac{3}{2})}\big),$$

and likewise

$$s_n(V_Y K_{0,Y}^* M_Y^*) = O\big(n^{-(H_Y+\frac{3}{2})}\big), \qquad s_n(V_Y V_Y^*) = s_n(V_Y)^2 = O\big(n^{-(2H_Y+3)}\big).$$

Because $H_Y < \frac{1}{2}$, all three exponents are strictly larger than $2H_Y + 1$. Hence

$$s_n(E_0) = o\big(n^{-(2H_Y+1)}\big),$$

so the remainder lies below the covariance scale. Since multiplication by boundedly invertible operators preserves spectral order in the min–max sense,

$$\lambda_n(M_Y C_{0,Y} M_Y^*) \asymp \lambda_n(C_{0,Y}).$$

For the pure fractional covariance,

$$\lambda_n(C_{0,Y}) \asymp s_n(K_{0,Y})^2 \asymp n^{-(2H_Y+1)}$$

by Birman–Solomyak. The Weyl–Ky Fan inequalities, applied to $C_Y = \nu_Y^2 M_Y C_{0,Y} M_Y^* + E_0$ and then to the reverse decomposition of the principal term, show that a compact perturbation with strictly smaller singular-value order cannot change the two-sided eigenvalue order. Therefore

$$\lambda_n(C_Y) \asymp n^{-(2H_Y+1)}.$$

□

Remark 4.2 (Transfer of spectral order). If K is a compact operator with

$$s_n(K) \asymp n^{-\alpha},$$

then the covariance operator $C = KK^*$ satisfies

$$\lambda_n(C) \asymp n^{-2\alpha}.$$

5. Pure fractional identification

Proposition 5.1 (Pure fractional benchmark). *Let $H_Y, H_{XY} \in (0, \frac{1}{2})$ with $H_{XY} > H_Y$, and let $K_{0,Y}, K_{0,XY}$ be the pure Riemann–Liouville operators from* (7)*. Define*

$$B_0 := (K_{0,Y}K_{0,Y}^*)^{-1/2}K_{0,XY}$$

initially on the natural domain of $(K_{0,Y}K_{0,Y}^*)^{-1/2}$. Then B_0 extends to a bounded operator on $L^2[0,T]$, and

$$s_n(B_0) \asymp n^{-(H_{XY}-H_Y)}.$$

Consequently,

$$s_n(B_0B_0^*) \asymp n^{-2(H_{XY}-H_Y)}.$$

Proof. Set

$$\delta := H_{XY} - H_Y > 0.$$

The semigroup property of Riemann–Liouville integrals gives

$$K_{0,XY} = c_\delta \, K_{0,Y} I_{0+}^\delta$$

for a nonzero constant c_δ. The Abel uniqueness theorem implies $\ker K_{0,Y} = \ker K_{0,Y}^* = \{0\}$. Hence the polar decomposition

$$K_{0,Y} = (K_{0,Y}K_{0,Y}^*)^{1/2}U_Y$$

has a unitary polar factor U_Y on $L^2[0,T]$. On the natural support of $(K_{0,Y}K_{0,Y}^*)^{-1/2}$,

$$(K_{0,Y}K_{0,Y}^*)^{-1/2}K_{0,Y} = U_Y.$$

Therefore

$$B_0 = c_\delta U_Y I_{0+}^\delta.$$

Unitary multiplication and the nonzero scalar c_δ do not change singular-value order. By Birman–Solomyak,

$$s_n(I_{0+}^\delta) \asymp n^{-\delta},$$

and hence

$$s_n(B_0) \asymp n^{-\delta}.$$

The second claim follows because the nonzero eigenvalues of $B_0B_0^*$ are the squared singular values of B_0.

□

6. Relative covariance operator and spectral structure

Spectral reduction. The relative covariance operator reduces to a Volterra factor whose kernel carries the roughness gap $H_{XY} - H_Y$. The case $H_{XY} > H_Y$ is the smoothing regime. In that regime the relative covariance operator is compact, and its spectral order is determined by the local fractional singularity.

The reduction has three levels. The operator A removes the covariance normalisation and isolates the cross-roughness mismatch. The decomposition

$$A = P_\delta + R_\delta$$

separates the principal fractional part from a lower-order weak-Schatten residual. Perturbation transfers the singular-value asymptotics from P_δ to A and from A to the relative covariance operator. The spectral exponent is the gap $H_{XY} - H_Y$.

6.1. **Preparatory kernel factorisations.**

Proposition 6.1 (Remainder factorisation through a fractional kernel)**.** *Let $\alpha > 0$ and $\rho > 0$. Let $R \in L^\infty(\Delta_T)$, and define*

$$(W_{\alpha,\rho,R}f)(t) := \int_0^t (t-s)^{\alpha+\rho-1} R(t,s) f(s)\, ds.$$

Assume that $W_{\alpha,\rho,R}$ maps $L^2[0,T]$ boundedly into $I_{0+}^{\alpha+\rho}(L^2[0,T])$. Then there exists a bounded operator $S_{\rho,R} : L^2[0,T] \to I_{0+}^{\rho}(L^2[0,T])$ such that

$$W_{\alpha,\rho,R} = I_{0+}^{\alpha} S_{\rho,R}. \tag{8}$$

Regarded as an operator on $L^2[0,T]$ through the canonical embedding $I_{0+}^{\rho}(L^2) \hookrightarrow L^2$, it satisfies

$$S_{\rho,R} \in \mathfrak{S}_p \qquad \text{for every } p > 1/\rho.$$

In particular, if $\rho > 1/q$, then $S_{\rho,R} \in \mathfrak{S}_q$.

Proof. The assumed mapping property places the range of $W_{\alpha,\rho,R}$ inside $I_{0+}^{\alpha+\rho}(L^2)$, hence inside $I_{0+}^{\alpha}(L^2)$. Define

$$S_{\rho,R} := D_{0+}^{\alpha} W_{\alpha,\rho,R}$$

on this range. Since

$$D_{0+}^{\alpha} : I_{0+}^{\alpha+\rho}(L^2) \longrightarrow I_{0+}^{\rho}(L^2)$$

is bounded, $S_{\rho,R}$ is bounded from $L^2[0,T]$ into $I_{0+}^{\rho}(L^2[0,T])$. The identity $I_{0+}^{\alpha} D_{0+}^{\alpha} u = u$ on $I_{0+}^{\alpha}(L^2)$ gives

$$I_{0+}^{\alpha} S_{\rho,R} = W_{\alpha,\rho,R},$$

which proves (8).

The Schatten order is the order of the residual embedding. Let

$$J_\rho : I_{0+}^{\rho}(L^2[0,T]) \hookrightarrow L^2[0,T]$$

be the canonical inclusion. By the Birman–Solomyak theorem for fractional integration operators,

$$s_n(J_\rho) \asymp n^{-\rho}.$$

The L^2-realisation of $S_{\rho,R}$ is $J_\rho S_{\rho,R}$, with $S_{\rho,R} : L^2 \to I_{0+}^{\rho}(L^2)$ bounded. Hence

$$s_n(S_{\rho,R}) = O(n^{-\rho}),$$

and $S_{\rho,R} \in \mathfrak{S}_p$ for every $p > 1/\rho$. The final assertion follows by taking $p = q$. ☐

Lemma 6.2 (Factorisation through the rougher kernel)**.** *Set*

$$\alpha_Y := H_Y + \tfrac{1}{2}, \qquad \delta := H_{XY} - H_Y > 0,$$

and define the diagonal traces

$$d_Y(s) := L_Y(s,s), \qquad d_{XY}(s) := L_{XY}(s,s).$$

Then there exist operators $S_Y, S_{XY} \in \mathcal{L}(L^2[0,T])$ such that

$$K_Y = K_{0,Y}(M_{d_Y} + S_Y), \tag{9}$$

and

$$K_{XY} = K_{0,Y}(c_\delta I_{0+}^{\delta} M_{d_{XY}} + S_{XY}), \tag{10}$$

where $c_\delta \neq 0$ is the constant from the semigroup identity

$$K_{0,XY} = c_\delta K_{0,Y} I_{0+}^{\delta}.$$

The residual factors satisfy:

(i) $M_{d_Y} + S_Y$ is boundedly invertible on $L^2[0,T]$;
(ii) $S_Y \in \mathfrak{S}_p$ for every $p > 1$;
(iii) $S_{XY} \in \mathfrak{S}_p$ for every $p > \frac{1}{\delta+1}$.

Proof. Because L_Y and L_{XY} are C^1 on a neighbourhood of Δ_T, the diagonal expansion gives

$$L_i(t,s) = L_i(s,s) + (t-s)R_i(t,s), \qquad i \in \{Y, XY\},$$

where

$$R_i(t,s) = \int_0^1 \partial_1 L_i(s + \theta(t-s), s)\, d\theta$$

is bounded on Δ_T. The standard fractional-Volterra mapping theorem places the resulting diagonal-vanishing C^1 remainders in the range required by Proposition 6.1; see [26, Chs. 2–3].

For K_Y, this gives

$$K_Y f(t) = \int_0^t (t-s)^{\alpha_Y - 1} d_Y(s) f(s)\, ds + \int_0^t (t-s)^{\alpha_Y} R_Y(t,s) f(s)\, ds.$$

Hence

$$K_Y = K_{0,Y} M_{d_Y} + W_Y,$$

where $W_Y = W_{\alpha_Y, 1, R_Y}$ in the notation of Proposition 6.1. That proposition, with $\alpha = \alpha_Y$ and $\rho = 1$, gives a bounded operator S_Y such that

$$W_Y = K_{0,Y} S_Y, \qquad S_Y \in \mathfrak{S}_p \quad \text{for every } p > 1.$$

Therefore

$$K_Y = K_{0,Y}(M_{d_Y} + S_Y).$$

Likewise,

$$K_{XY} = K_{0,XY} M_{d_{XY}} + W_{XY} = c_\delta K_{0,Y} I^\delta_{0+} M_{d_{XY}} + W_{XY},$$

where $W_{XY} = W_{\alpha_Y, \delta+1, R_{XY}}$. Proposition 6.1, with $\alpha = \alpha_Y$ and $\rho = \delta + 1$, gives a bounded operator S_{XY} such that

$$W_{XY} = K_{0,Y} S_{XY}, \qquad S_{XY} \in \mathfrak{S}_p \quad \text{for every } p > \frac{1}{\delta+1}.$$

Hence

$$K_{XY} = K_{0,Y}(c_\delta I^\delta_{0+} M_{d_{XY}} + S_{XY}).$$

Since $d_Y(s) = L_Y(s,s) \geq c_Y > 0$, the multiplication operator M_{d_Y} is boundedly invertible. Because S_Y is compact,

$$M_{d_Y} + S_Y = M_{d_Y}(I + M_{d_Y}^{-1} S_Y)$$

is Fredholm of index 0. It is injective. Indeed, if $(M_{d_Y} + S_Y)f = 0$, then

$$K_Y f = K_{0,Y}(M_{d_Y} + S_Y) f = 0.$$

The Abel–Volterra uniqueness theorem for kernels with nonvanishing diagonal amplitude gives $f = 0$; see [26, Chs. 2–3]. Fredholm index zero then gives surjectivity and bounded invertibility. □

6.2. **Weak Schatten notation.** For $p > 0$, let $\mathfrak{S}_{p,\infty}$ denote the weak Schatten ideal (see Simon [27] for background on trace ideals and Schatten classes):

$$\mathfrak{S}_{p,\infty} := \Big\{ K \in \mathcal{K}(L^2[0,T]) : s_n(K) = O(n^{-1/p}) \Big\}.$$

We write $(\mathfrak{S}_{p,\infty})_0$ for the closure of finite-rank operators in $\mathfrak{S}_{p,\infty}$, equivalently

$$(\mathfrak{S}_{p,\infty})_0 = \Big\{ K \in \mathcal{K}(L^2[0,T]) : s_n(K) = o(n^{-1/p}) \Big\}.$$

Lemma 6.3 (Weak-Schatten multiplication and perturbation)**.** *Let $p, q > 0$ and $r > 0$ satisfy*

$$\frac{1}{r} = \frac{1}{p} + \frac{1}{q}.$$

Then the following hold.

(i) If $A \in \mathfrak{S}_p$ and $B \in \mathfrak{S}_{q,\infty}$, then

$$AB,\ BA \in \mathfrak{S}_{r,\infty}.$$

(ii) If $A \in \mathfrak{S}_{p,\infty}$ and $B \in (\mathfrak{S}_{p,\infty})_0$, then

$$s_n(A+B) \asymp s_n(A)$$

whenever $s_n(A) \asymp n^{-1/p}$.

(iii) More explicitly, if

$$s_n(A) \asymp n^{-\alpha}, \qquad s_n(B) = o(n^{-\alpha}),$$

then

$$s_n(A+B) \asymp n^{-\alpha}.$$

Proof. Part (i) is the standard weak-Schatten Hölder rule. For the perturbation statement, use Ky Fan's inequality

$$s_{m+n-1}(X+Y) \leq s_m(X) + s_n(Y).$$

If $s_n(A) \asymp n^{-\alpha}$ and $s_n(B) = o(n^{-\alpha})$, then the upper bound follows from Ky Fan applied to $A+B$. For the lower bound, write $A = (A+B) - B$ and apply the same inequality:

$$s_{2n-1}(A) \leq s_n(A+B) + s_n(B).$$

Since $s_n(B) = o(n^{-\alpha})$ and $s_{2n-1}(A) \asymp n^{-\alpha}$, this gives $s_n(A+B) \geq cn^{-\alpha}$ for all sufficiently large n. This proves (iii). Part (ii) is the case $\alpha = 1/p$, because $B \in (\mathfrak{S}_{p,\infty})_0$ is exactly $s_n(B) = o(n^{-1/p})$. □

6.3. **Factorisation of covariance operators.** Recall from (6) that

$$C_Y = \nu_Y^2 K_Y K_Y^*, \qquad C_{XY} = c^2 \nu_Y^2 K_{XY} K_{XY}^*.$$

Since multiplication of a positive covariance operator by a strictly positive scalar does not affect its Cameron–Martin space, relative covariance structure, or Schatten membership, it is convenient to work with the normalised covariance operators

$$\widehat{C}_Y := K_Y K_Y^*, \qquad \widehat{C}_{XY} := K_{XY} K_{XY}^*.$$

Then

$$C_Y = \nu_Y^2 \widehat{C}_Y, \qquad C_{XY} = c^2 \nu_Y^2 \widehat{C}_{XY}.$$

Recall also from Lemma 6.2 that

$$K_Y = K_{0,Y} A_Y, \qquad K_{XY} = K_{0,Y} A_{XY},$$

where

$$A_Y = M_{d_Y} + S_Y, \qquad A_{XY} = c_\delta I_{0+}^\delta M_{d_{XY}} + S_{XY}.$$

Define

$$A := A_Y^{-1} A_{XY}.$$

Then

$$K_{XY} = K_Y A, \qquad \widehat{C}_{XY} = K_Y A A^* K_Y^*.$$

Operator domains. All inverse and fractional-power operators below are interpreted by spectral calculus on the natural support spaces of the corresponding covariance operators. Operators such as $C_Y^{-1/2}$ act on $\overline{\operatorname{Ran}(C_Y^{1/2})}$, and identities involving relative covariance operators are understood on that space.

6.4. **Reduction to the relative operator.** Define the normalised relative factor

$$\widehat{B} := \widehat{C}_Y^{-1/2}\widehat{C}_{XY}^{1/2}.$$

The Volterra uniqueness theorem gives $\ker K_Y = \ker K_Y^* = \ker K_{XY} = \ker K_{XY}^* = \{0\}$. The polar factors in the decompositions

$$K_Y = \widehat{C}_Y^{1/2}U_Y, \qquad K_{XY} = \widehat{C}_{XY}^{1/2}U_{XY}$$

are therefore unitary on $L^2[0,T]$. Using $K_{XY} = K_Y A$, we obtain

$$\widehat{C}_{XY}^{1/2}U_{XY} = \widehat{C}_Y^{1/2}U_Y A,$$

hence

$$\widehat{B}\, U_{XY} = U_Y A.$$

Lemma 6.4 (Singular value equivalence)**.** *The operators $\widehat{B}$ and A have the same nonzero singular values:*

$$s_n(\widehat{B}) = s_n(A).$$

Proof. From $\widehat{B}\, U_{XY} = U_Y A$ and the unitarity of the two polar factors,

$$\widehat{B} = U_Y A U_{XY}^*.$$

Unitary equivalence preserves the singular spectrum. Hence $\widehat{B}$ and A have the same nonzero singular values, with the same multiplicities. □

Therefore the normalised relative covariance operator

$$\widehat{\mathcal{R}}_{Y,XY} := \widehat{C}_Y^{-1/2}\widehat{C}_{XY}\widehat{C}_Y^{-1/2} = \widehat{B}\,\widehat{B}^*$$

has eigenvalues

$$\lambda_n(\widehat{\mathcal{R}}_{Y,XY}) = s_n(A)^2.$$

Thus the spectral analysis reduces to A.

6.5. **Decomposition of the relative operator.** Recall from Lemma 6.2 that

$$A_Y = M_{d_Y} + S_Y, \qquad A_{XY} = c_\delta I_{0+}^{\delta} M_{d_{XY}} + S_{XY},$$

with

$$\delta := H_{XY} - H_Y.$$

Assume $\delta > 0$.

Define the principal operator

$$P_\delta := c_\delta M_{d_Y}^{-1} I_{0+}^{\delta} M_{d_{XY}}.$$

Proposition 6.5 (Principal + remainder decomposition)**.** *The operator $A = A_Y^{-1}A_{XY}$ admits the decomposition*

$$A = P_\delta + R_\delta,$$

where R_δ is compact and satisfies

$$s_n(R_\delta) = o(n^{-\delta}).$$

Proof. Write

$$A - P_\delta = c_\delta(A_Y^{-1} - M_{d_Y}^{-1})I_{0+}^{\delta} M_{d_{XY}} + A_Y^{-1}S_{XY}.$$

Each term is lower order.

Step 1: expansion of A_Y^{-1}. Since

$$A_Y = M_{d_Y}(I + M_{d_Y}^{-1} S_Y),$$

inversion gives

$$A_Y^{-1} = (I + M_{d_Y}^{-1} S_Y)^{-1} M_{d_Y}^{-1}.$$

Hence

$$A_Y^{-1} - M_{d_Y}^{-1} = -M_{d_Y}^{-1} S_Y A_Y^{-1}.$$

Since $S_Y \in \mathfrak{S}_p$ for every $p > 1$, we have

$$A_Y^{-1} - M_{d_Y}^{-1} \in \mathfrak{S}_p \quad \text{for all } p > 1.$$

Step 2: composition with I_{0+}^{δ}. It is classical that

$$s_n(I_{0+}^{\delta}) \asymp n^{-\delta}.$$

See, for example, Dostanić [22], Tuan and Gorenflo [23], and Gorenflo and Tuan [24]. Hence

$$I_{0+}^{\delta} \in \mathfrak{S}_{1/\delta,\infty}.$$

Thus, for any $p > 1$,

$$(A_Y^{-1} - M_{d_Y}^{-1}) I_{0+}^{\delta} \in \mathfrak{S}_{r,\infty}, \quad \frac{1}{r} = \frac{1}{p} + \delta.$$

Since $1/r = \delta + 1/p$ is strictly larger than δ, this gives the stronger decay

$$s_n((A_Y^{-1} - M_{d_Y}^{-1}) I_{0+}^{\delta}) = O(n^{-1/r}) = o(n^{-\delta}).$$

Multiplication by bounded operators preserves this estimate.
Step 3: control of $A_Y^{-1} S_{XY}$. Since A_Y^{-1} is bounded and

$$S_{XY} \in \mathfrak{S}_p \quad \text{for every } p > \frac{1}{\delta + 1},$$

we have

$$s_n(A_Y^{-1} S_{XY}) = O(n^{-1/p}).$$

The interval

$$\left(\frac{1}{\delta + 1}, \frac{1}{\delta} \right)$$

is nonempty. Choosing p in this interval gives $1/p > \delta$, and therefore

$$s_n(A_Y^{-1} S_{XY}) = o(n^{-\delta}).$$

Conclusion. Combining both terms gives

$$s_n(R_\delta) = o(n^{-\delta}).$$

□

6.6. **Singular value asymptotics.**

Proposition 6.6 (Weighted fractional operator)**.** *Assume d_Y, d_{XY} are bounded above and below by positive constants. Then*

$$s_n(P_\delta) \asymp n^{-\delta}.$$

Proof. Write

$$P_\delta = c_\delta M_{d_Y}^{-1} I_{0+}^{\delta} M_{d_{XY}}.$$

The multipliers $M_{d_Y}^{-1}$ and $M_{d_{XY}}$ are bounded and boundedly invertible, and $c_\delta \neq 0$. Hence

$$s_n(P_\delta) \leq |c_\delta| \, \|M_{d_Y}^{-1}\| \, \|M_{d_{XY}}\| \, s_n(I_{0+}^{\delta}),$$

while the reverse representation of I_{0+}^{δ} in terms of P_δ gives the opposite bound up to a constant. Since $s_n(I_{0+}^{\delta}) \asymp n^{-\delta}$, the claim follows. □

Theorem 6.7 (Spectral asymptotics of A)**.** *Assume $\delta > 0$. Then*

$$s_n(A) \asymp n^{-\delta}.$$

Proof. From Proposition 6.5,

$$A = P_\delta + R_\delta, \quad s_n(R_\delta) = o(n^{-\delta}),$$

and from Proposition 6.6,

$$s_n(P_\delta) \asymp n^{-\delta}.$$

The result follows from Lemma 6.3(iii). ☐

6.7. Spectral structure of the covariance operator.

Theorem 6.8 (Eigenvalue asymptotics of the relative covariance operator)**.** *Let* $\delta = H_{XY} - H_Y > 0$*. Then:*

(i)

$$s_n(\widehat{B}) \asymp n^{-\delta}.$$

(ii)

$$\lambda_n(\widehat{\mathcal{R}}_{Y,XY}) = s_n(\widehat{B})^2 \asymp n^{-2\delta}.$$

(iii)

$$\widehat{\mathcal{R}}_{Y,XY} \in \mathfrak{S}_2 \quad \Longleftrightarrow \quad \delta > \frac{1}{4}.$$

Proof. By Theorem 6.7 and Lemma 6.4,

$$s_n(\widehat{B}) = s_n(A) \asymp n^{-\delta},$$

which proves (i). Since $\widehat{\mathcal{R}}_{Y,XY} = \widehat{B}\,\widehat{B}^*$, the nonzero eigenvalues of $\widehat{\mathcal{R}}_{Y,XY}$ are the squares of the singular values of $\widehat{B}$. Thus

$$\lambda_n(\widehat{\mathcal{R}}_{Y,XY}) = s_n(\widehat{B})^2 \asymp n^{-2\delta},$$

which proves (ii). Finally,

$$\widehat{\mathcal{R}}_{Y,XY} \in \mathfrak{S}_2 \iff \sum_{n\geq 1} \lambda_n(\widehat{\mathcal{R}}_{Y,XY})^2 < \infty \iff \sum_{n\geq 1} n^{-4\delta} < \infty \iff \delta > \frac{1}{4}.$$

This proves (iii). ☐

Remark 6.9 (Role of weak Schatten bounds). The weak-Schatten estimates are used only to place the residual terms below the principal fractional order. The two-sided order comes from the weighted fractional operator P_δ and is then preserved by the smaller remainder.

Since $C_Y = \nu_Y^2 \widehat{C}_Y$ and $C_{XY} = c^2 \nu_Y^2 \widehat{C}_{XY}$, the relative covariance operator with respect to the original (unnormalised) covariances is

$$\mathcal{R}_{Y,XY} := C_Y^{-1/2} C_{XY} C_Y^{-1/2} = c^2 \widehat{\mathcal{R}}_{Y,XY}.$$

If $c > 0$, then $\lambda_n(\mathcal{R}_{Y,XY}) = c^2 \lambda_n(\widehat{\mathcal{R}}_{Y,XY}) \asymp n^{-2\delta}$ and $\mathcal{R}_{Y,XY} \in \mathfrak{S}_2$ if and only if $\widehat{\mathcal{R}}_{Y,XY} \in \mathfrak{S}_2$. If $c = 0$, then $\mathcal{R}_{Y,XY} = 0$ and the coupled law coincides with the uncoupled law.

6.8. Cameron–Martin spaces.

Lemma 6.10 (Cameron–Martin spaces under bounded relative perturbations)**.** *Let C be a positive injective trace-class operator on a separable Hilbert space $\mathcal{H}$, and let T be a bounded positive self-adjoint operator on* $\overline{\mathrm{Ran}(C^{1/2})}$*. Define*

$$\widetilde{C} := C^{1/2}(I + T)C^{1/2}.$$

Then

$$\mathrm{Ran}(\widetilde{C}^{1/2}) = \mathrm{Ran}(C^{1/2}),$$

and the Cameron–Martin norms induced by C and $\widetilde{C}$ are equivalent.

Proof. Let $A := (I+T)^{1/2}$. Since T is bounded and positive, A is bounded and boundedly invertible on $\overline{\mathrm{Ran}(C^{1/2})}$. In addition,

$$\widetilde{C} = C^{1/2}A^2C^{1/2} = (C^{1/2}A)(C^{1/2}A)^*.$$

By the Douglas range identity,

$$\mathrm{Ran}(\widetilde{C}^{1/2}) = \mathrm{Ran}(C^{1/2}A).$$

Because A is a bounded isomorphism of $\overline{\mathrm{Ran}(C^{1/2})}$,

$$\mathrm{Ran}(C^{1/2}A) = \mathrm{Ran}(C^{1/2}).$$

The Cameron–Martin norm is the quotient norm induced by the corresponding covariance square root. The bounded invertibility of A on the support space gives equivalence of these quotient norms. □

6.9. **Feldman–Hájek threshold.**

Theorem 6.11 (Equivalence threshold in the smoothing regime)**.** *Let $\delta := H_{XY} - H_Y > 0$ and assume $c > 0$. Then the Gaussian measures induced by C_Y and $C_Y + C_{XY}$ are equivalent if and only if*

$$\delta > \frac{1}{4}.$$

They are mutually singular when $0 < \delta \leq \frac{1}{4}$.

Remark 6.12 (Non-smoothing regime). The cases $\delta \leq 0$ are not covered by Theorem 6.11. In the pure fractional setting, the behaviour is consistent with mutual singularity when $H_{XY} \leq H_Y$, but the corresponding statement for the present modulated model requires a separate argument and is not used here.

Proof. All operators are considered on $\overline{\mathrm{Ran}(C_Y^{1/2})}$, where the functional calculus is defined. Since $C_Y = \nu_Y^2 \widehat{C}_Y$ and $C_{XY} = c^2\nu_Y^2\widehat{C}_{XY}$, we have

$$C_Y + C_{XY} = C_Y^{1/2}(I + c^2\widehat{\mathcal{R}}_{Y,XY})C_Y^{1/2} = C_Y^{1/2}(I + \mathcal{R}_{Y,XY})C_Y^{1/2},$$

where $\mathcal{R}_{Y,XY} = c^2\widehat{\mathcal{R}}_{Y,XY}$. Since $\mathcal{R}_{Y,XY}$ is bounded and positive, Lemma 6.10 implies that the Cameron–Martin spaces of C_Y and $C_Y + C_{XY}$ coincide.

The Cameron–Martin condition in Feldman–Hájek has therefore already been verified. The remaining condition is the Hilbert–Schmidt condition for the relative square-root perturbation. By the Feldman–Hájek theorem on $L^2[0,T]$, the two centred Gaussian measures with covariance operators C_Y and $C_Y + C_{XY}$ are equivalent if and only if

$$(I + \mathcal{R}_{Y,XY})^{1/2} - I \in \mathfrak{S}_2.$$

Because $\mathcal{R}_{Y,XY}$ is positive compact, its eigenvalues tend to zero and functional calculus gives

$$(\sqrt{1+\lambda} - 1)^2 \asymp \lambda^2 \qquad (\lambda \downarrow 0),$$

so

$$(I + \mathcal{R}_{Y,XY})^{1/2} - I \in \mathfrak{S}_2 \quad \Longleftrightarrow \quad \mathcal{R}_{Y,XY} \in \mathfrak{S}_2.$$

Because multiplication by the nonzero scalar c^2 does not affect Hilbert–Schmidt membership, this reduces to $\widehat{\mathcal{R}}_{Y,XY} \in \mathfrak{S}_2$. By Theorem 6.8(iii),

$$\widehat{\mathcal{R}}_{Y,XY} \in \mathfrak{S}_2 \quad \Longleftrightarrow \quad \delta > \tfrac{1}{4}.$$

When $0 < \delta \leq \frac{1}{4}$, the Cameron–Martin spaces coincide but the Hilbert–Schmidt condition fails; the Feldman–Hájek dichotomy therefore gives mutual singularity. Applying Proposition 2.6 yields the same dichotomy for the path laws μ_η and μ_0 on $C([0,T])$. □

Corollary 6.13 (Sharp latent-mode growth). *Assume $\delta > 0$ and $c > 0$, and define*

$$\mathcal{E}_{\mathrm{latent}}(N) := \sum_{n=1}^{N} \lambda_n(\mathcal{R}_{Y,XY}).$$

Then

$$\mathcal{E}_{\mathrm{latent}}(N) \asymp \begin{cases} N^{1-2\delta}, & 0 < \delta < \frac{1}{2}, \\ \log N, & \delta = \frac{1}{2}, \\ 1, & \delta > \frac{1}{2}. \end{cases}$$

Under the standing Hurst range $H_Y, H_{XY} \in (0, \frac{1}{2})$, one has $\delta < \frac{1}{2}$ whenever $\delta > 0$; the logarithmic and bounded alternatives record the general summability calculus for the same spectral sequence. In particular, in the latent regime $0 < \delta \le \frac{1}{4}$, the cumulative spectral mass diverges polynomially.

Proof. This is the partial-sum asymptotic associated with $\lambda_n(\mathcal{R}_{Y,XY}) \asymp n^{-2\delta}$. □

Corollary 6.14 (Finite-resolution latent energy). *Assume $c > 0$. Let $\Delta \in (0, 1)$ be an observation scale and set $N_\Delta := \lfloor \Delta^{-1} \rfloor$. If $0 < \delta < \frac{1}{2}$, then*

$$\mathcal{E}_{\mathrm{latent}}(N_\Delta) \asymp \Delta^{-(1-2\delta)}.$$

Proof. Immediate from Corollary 6.13. □

Remark 6.15 (On the regime $H_{XY} \le H_Y$). Theorem 6.11 treats the smoothing regime $H_{XY} > H_Y$. In the pure fractional model, the relative factor B_0 ceases to be compact when $H_{XY} = H_Y$ and is no longer bounded when $H_{XY} < H_Y$, which provides heuristic evidence for mutual singularity. The analogous statement for the modulated setting should be formulated separately, rather than folded into Theorem 6.11 without proof.

7. Short-maturity skew under the pricing measure

Structure under the pricing measure. The Bachelier specification of Section 2 fixes an additive price process. The short-maturity skew is the *Bachelier* implied-volatility skew. The Gaussian volatility field is kept unchanged under $\mathbb{P}$ and $\mathbb{Q}$; the pricing interpretation changes, not the second-order Volterra geometry.

In the Bachelier setting, the leading nonzero contribution to the third cumulant of the return is the leverage term

$$\mathbb{E}\left[B_T^2 \int_0^T \mathcal{Y}_t \, dB_t\right],$$

not a Gaussian odd moment. This term produces the rough-volatility scaling $T^{H-\frac{1}{2}}$ for the implied-volatility skew and yields the pricing-level threshold $H_{XY} = H_Y$.

7.1. Model and asymptotic regime.

The notation of Section 2 is retained. Let

$$V_t = \nu_Y \int_0^t (t-s)^{H_Y - \frac{1}{2}} L_Y(t,s) \, dW_s^\perp, \qquad Z_t = \int_0^t (t-s)^{H_{XY} - \frac{1}{2}} L_{XY}(t,s) \, dW_s^X,$$

with $H_Y, H_{XY} \in (0, \frac{1}{2})$ and with L_Y, L_{XY} continuous and diagonally nondegenerate:

$$0 < c_- \le L_Y(s,s), \, L_{XY}(s,s) \le c_+.$$

Set

$$\mathcal{Y}_t := V_t + c\nu_Y Z_t - C(t), \qquad \sigma_t := \sigma_0 e^{\mathcal{Y}_t},$$

where $C(t)$ is the deterministic convexity correction from Section 2. Let

$$B_t := \rho W_t^X + \sqrt{1-\rho^2} \, W_t^\perp, \qquad |\rho| < 1,$$

and consider the Bachelier price dynamics under $\mathbb{Q}$

$$dS_t = \sigma_t \, dB_t.$$

The short-maturity object is the return

$$R_T := S_T - S_0 = \int_0^T \sigma_t \, dB_t$$

as $T \downarrow 0$.

Unless explicitly stated otherwise, all expectations and probabilities in this pricing section are taken under $\mathbb{Q}$.

Since the model is additive, $\mathbb{E}[R_T] = 0$, and therefore the third cumulant coincides with the third moment. We write

$$\kappa_3(T) := \mathbb{E}[R_T^3].$$

Set

$$H_* := \min(H_Y, H_{XY}).$$

Lemma 7.1 (Uniform Gaussian moment scale)**.** *For each fixed $m \geq 1$,*

$$\sup_{t \leq T} \mathbb{E}[|\mathcal{Y}_t|^m] = O(T^{mH_*}) \qquad \text{as } T \downarrow 0. \tag{11}$$

Proof. By Itô isometry and the boundedness of L_Y and L_{XY},

$$\mathbb{E}[V_t^2] = \nu_Y^2 \int_0^t (t-s)^{2H_Y - 1} L_Y(t,s)^2 \, ds = O(t^{2H_Y}),$$

and similarly

$$\mathbb{E}[Z_t^2] = O(t^{2H_{XY}}).$$

Hence

$$\mathrm{Var}(V_t + c\nu_Y Z_t) = O(t^{2H_*})$$

uniformly for $t \leq T$. The convexity correction $C(t)$ is deterministic and satisfies $C(t) = O(t^{2H_*})$, so

$$\mathbb{E}[\mathcal{Y}_t^2] = O(t^{2H_*}).$$

Since $\mathcal{Y}_t$ is Gaussian up to the deterministic shift $C(t)$, standard Gaussian moment equivalence yields

$$\mathbb{E}[|\mathcal{Y}_t|^m] \leq C_m \big(\mathbb{E}[\mathcal{Y}_t^2]^{m/2} + |C(t)|^m\big) = O(t^{mH_*}),$$

and the stated uniform bound follows by taking the supremum over $t \leq T$. $\square$

7.2. **Leverage covariance asymptotics.** The leading skew contribution is governed by the covariance between the price Brownian motion B_t and the volatility factor $\mathcal{Y}_t$.

Lemma 7.2 (Leverage covariance asymptotics)**.** *As $t \downarrow 0$,*

$$\mathbb{E}[B_t V_t] = \frac{\nu_Y \sqrt{1-\rho^2}\, L_Y(0,0)}{H_Y + \frac{1}{2}} \, t^{H_Y + \frac{1}{2}} + o\Big(t^{H_Y + \frac{1}{2}}\Big),$$

and

$$\mathbb{E}[B_t \, c\nu_Y Z_t] = \frac{c\nu_Y \rho \, L_{XY}(0,0)}{H_{XY} + \frac{1}{2}} \, t^{H_{XY} + \frac{1}{2}} + o\Big(t^{H_{XY} + \frac{1}{2}}\Big).$$

Consequently,

$$\mathbb{E}[B_t \mathcal{Y}_t] = a_Y \, t^{H_Y + \frac{1}{2}} + a_{XY} \, t^{H_{XY} + \frac{1}{2}} + o\Big(t^{H_* + \frac{1}{2}}\Big), \tag{12}$$

where

$$a_Y := \frac{\nu_Y \sqrt{1-\rho^2}\, L_Y(0,0)}{H_Y + \frac{1}{2}}, \qquad a_{XY} := \frac{c\nu_Y \rho \, L_{XY}(0,0)}{H_{XY} + \frac{1}{2}}.$$

In particular, $a_Y > 0$, while a_{XY} has the sign of ρc and may vanish if $\rho c = 0$.

Proof. Since $B_t = \rho W_t^X + \sqrt{1-\rho^2}\, W_t^\perp$ and $W^X, W^\perp$ are independent,

$$\mathbb{E}[B_t V_t] = \sqrt{1-\rho^2}\, \nu_Y \int_0^t (t-s)^{H_Y - \frac{1}{2}} L_Y(t,s)\, ds.$$

Rescale $s = tu$ to obtain

$$\mathbb{E}[B_t V_t] = \sqrt{1-\rho^2}\, \nu_Y\, t^{H_Y + \frac{1}{2}} \int_0^1 (1-u)^{H_Y - \frac{1}{2}} L_Y(t,tu)\, du.$$

Since L_Y is continuous on Δ_T, $L_Y(t,tu) \to L_Y(0,0)$ uniformly in $u \in [0,1]$ as $t \downarrow 0$. By dominated convergence,

$$\int_0^1 (1-u)^{H_Y - \frac{1}{2}} L_Y(t,tu)\, du \to L_Y(0,0) \int_0^1 (1-u)^{H_Y - \frac{1}{2}}\, du = \frac{L_Y(0,0)}{H_Y + \frac{1}{2}},$$

which proves the first asymptotic formula.

Similarly,

$$\mathbb{E}[B_t\, c\nu_Y Z_t] = c\nu_Y \rho \int_0^t (t-s)^{H_{XY} - \frac{1}{2}} L_{XY}(t,s)\, ds = c\nu_Y \rho\, t^{H_{XY} + \frac{1}{2}} \int_0^1 (1-u)^{H_{XY} - \frac{1}{2}} L_{XY}(t,tu)\, du,$$

and the same argument gives the second asymptotic formula.

Finally, $\mathbb{E}[B_t C(t)] = C(t)\mathbb{E}[B_t] = 0$, so

$$\mathbb{E}[B_t \mathcal{Y}_t] = \mathbb{E}[B_t V_t] + \mathbb{E}[B_t\, c\nu_Y Z_t],$$

which gives (12). □

7.3. **Linearised return and the third cumulant.**

Lemma 7.3 (Linearised return decomposition)**.** *As $T \downarrow 0$,*

$$R_T = \sigma_0 B_T + \sigma_0 \int_0^T \mathcal{Y}_t\, dB_t + \mathcal{R}_T,$$

where

$$\mathbb{E}\Big[|\mathcal{R}_T|^3\Big] = o\Big(T^{H_* + \frac{3}{2}}\Big).$$

More generally, for each fixed $p \geq 2$,

$$\mathbb{E}[|\mathcal{R}_T|^p] = O\Big(T^{2pH_* + \frac{p}{2}}\Big).$$

Proof. Write

$$e^{\mathcal{Y}_t} = 1 + \mathcal{Y}_t + q(\mathcal{Y}_t), \qquad q(y) := e^y - 1 - y.$$

Then

$$R_T = \sigma_0 \int_0^T dB_t + \sigma_0 \int_0^T \mathcal{Y}_t\, dB_t + \sigma_0 \int_0^T q(\mathcal{Y}_t)\, dB_t,$$

so it suffices to set

$$\mathcal{R}_T := \sigma_0 \int_0^T q(\mathcal{Y}_t)\, dB_t.$$

Since $|q(y)| \leq C y^2 e^{C|y|}$, the Gaussian moment bound (11) implies

$$\sup_{t \leq T} \mathbb{E}[|q(\mathcal{Y}_t)|^m] = O(T^{2mH_*}) \qquad (m \geq 1).$$

For each fixed $p \geq 2$, Burkholder–Davis–Gundy and Jensen's inequality in time give

$$\mathbb{E}[|\mathcal{R}_T|^p] \leq C_p\, \mathbb{E}\left[\left(\int_0^T |q(\mathcal{Y}_t)|^2\, dt\right)^{p/2}\right] \leq C_p\, T^{\frac{p}{2}-1} \int_0^T \mathbb{E}[|q(\mathcal{Y}_t)|^p]\, dt.$$

Using the moment estimate above,

$$\mathbb{E}[|\mathcal{R}_T|^p] = O\Big(T^{2pH_* + \frac{p}{2}}\Big).$$

For $p = 3$, this becomes

$$\mathbb{E}[|\mathcal{R}_T|^3] = O\Big(T^{6H_* + \frac{3}{2}}\Big) = o\Big(T^{H_* + \frac{3}{2}}\Big),$$

since $H_* > 0$. □

Proposition 7.4 (Leading cumulant contribution)**.** *As $T \downarrow 0$,*

$$\kappa_3(T) = 6\sigma_0^3 \int_0^T \mathbb{E}[B_t \mathcal{Y}_t]\, dt + o\Big(T^{H_*+\frac{3}{2}}\Big). \tag{13}$$

Proof. Set

$$G_T := \sigma_0 B_T, \qquad H_T := \sigma_0 \int_0^T \mathcal{Y}_t\, dB_t.$$

By Lemma 7.3,

$$R_T = G_T + H_T + \mathcal{R}_T.$$

Since G_T is centered Gaussian, $\mathbb{E}[G_T^3] = 0$. Expanding the cube and using Hölder's inequality, it is enough to show that every term other than $3\mathbb{E}[G_T^2 H_T]$ is $o(T^{H_*+\frac{3}{2}})$.

For each fixed $m \geq 1$, Burkholder–Davis–Gundy and Lemma 7.1 give

$$\mathbb{E}[|H_T|^m] \leq C_m\, \mathbb{E}\left[\left(\int_0^T |\mathcal{Y}_t|^2\, dt\right)^{m/2}\right] = O\Big(T^{mH_*+\frac{m}{2}}\Big).$$

First, by BDG and (11),

$$\mathbb{E}[|H_T|^3] \leq C\, \mathbb{E}\left[\left(\int_0^T |\mathcal{Y}_t|^2\, dt\right)^{3/2}\right] = O\Big(T^{3H_*+\frac{3}{2}}\Big) = o\Big(T^{H_*+\frac{3}{2}}\Big).$$

Also,

$$\mathbb{E}[|G_T H_T^2|] \leq \mathbb{E}[|G_T|^3]^{1/3}\mathbb{E}[|H_T|^3]^{2/3} = O(T^{1/2})\, O(T^{2H_*+1}) = O(T^{2H_*+\frac{3}{2}}) = o(T^{H_*+\frac{3}{2}}).$$

$$\mathbb{E}[|G_T|^m] = O(T^{m/2}) \qquad (m \geq 1),$$

because $G_T = \sigma_0 B_T$ is Gaussian. The terms involving $\mathcal{R}_T$ are therefore controlled explicitly:

$$\mathbb{E}[|G_T^2 \mathcal{R}_T|] \leq \mathbb{E}[|G_T|^4]^{1/2}\mathbb{E}[|\mathcal{R}_T|^2]^{1/2} = O(T)\, O\Big(T^{2H_*+\frac{1}{2}}\Big) = O\Big(T^{2H_*+\frac{3}{2}}\Big),$$

$$\mathbb{E}[|G_T \mathcal{R}_T^2|] \leq \mathbb{E}[|G_T|^2]^{1/2}\mathbb{E}[|\mathcal{R}_T|^4]^{1/2} = O\Big(T^{\frac{1}{2}}\Big)\, O\Big(T^{4H_*+1}\Big) = O\Big(T^{4H_*+\frac{3}{2}}\Big),$$

$$\mathbb{E}[|H_T^2 \mathcal{R}_T|] \leq \mathbb{E}[|H_T|^4]^{1/2}\mathbb{E}[|\mathcal{R}_T|^2]^{1/2} = O\Big(T^{2H_*+1}\Big)\, O\Big(T^{2H_*+\frac{1}{2}}\Big) = O\Big(T^{4H_*+\frac{3}{2}}\Big),$$

$$\mathbb{E}[|H_T \mathcal{R}_T^2|] \leq \mathbb{E}[|H_T|^2]^{1/2}\mathbb{E}[|\mathcal{R}_T|^4]^{1/2} = O\Big(T^{H_*+\frac{1}{2}}\Big)\, O\Big(T^{4H_*+1}\Big) = O\Big(T^{5H_*+\frac{3}{2}}\Big),$$

and

$$\mathbb{E}[|G_T H_T \mathcal{R}_T|] \leq \mathbb{E}[|G_T|^3]^{1/3}\mathbb{E}[|H_T|^3]^{1/3}\mathbb{E}[|\mathcal{R}_T|^3]^{1/3} = o\Big(T^{\frac{4H_*}{3}+\frac{3}{2}}\Big) = o\Big(T^{H_*+\frac{3}{2}}\Big).$$

The remaining pure remainder term satisfies

$$\mathbb{E}[|\mathcal{R}_T|^3] = o\Big(T^{H_*+\frac{3}{2}}\Big)$$

by Lemma 7.3. Therefore

$$\kappa_3(T) = 3\, \mathbb{E}[G_T^2 H_T] + o(T^{H_*+\frac{3}{2}}).$$

The surviving term is $\mathbb{E}[G_T^2 H_T]$. Since

$$B_T^2 = T + 2\int_0^T B_t\, dB_t,$$

we have

$$\mathbb{E}\left[B_T^2 \int_0^T \mathcal{Y}_t\, dB_t\right] = 2\, \mathbb{E}\left[\left(\int_0^T B_t\, dB_t\right)\left(\int_0^T \mathcal{Y}_t\, dB_t\right)\right],$$

because $\mathbb{E}[\int_0^T \mathcal{Y}_t\, dB_t] = 0$. By Itô isometry for stochastic integrals,

$$\mathbb{E}\left[\left(\int_0^T B_t\, dB_t\right)\left(\int_0^T \mathcal{Y}_t\, dB_t\right)\right] = \int_0^T \mathbb{E}[B_t \mathcal{Y}_t]\, dt.$$

Hence

$$\mathbb{E}[G_T^2 H_T] = 2\sigma_0^3 \int_0^T \mathbb{E}[B_t \mathcal{Y}_t]\, dt,$$

and (13) follows. □

Theorem 7.5 (Sharp small-time order of the third cumulant)**.** *As $T \downarrow 0$,*

$$\kappa_3(T) = C_Y^\kappa \, T^{H_Y+\frac{3}{2}} + C_{XY}^\kappa \, T^{H_{XY}+\frac{3}{2}} + o\Big(T^{H_*+\frac{3}{2}}\Big),$$

where

$$C_Y^\kappa := \frac{6\sigma_0^3 \nu_Y \sqrt{1-\rho^2}\, L_Y(0,0)}{\left(H_Y+\frac{1}{2}\right)\left(H_Y+\frac{3}{2}\right)},$$

and

$$C_{XY}^\kappa := \frac{6\sigma_0^3 c\nu_Y \rho\, L_{XY}(0,0)}{\left(H_{XY}+\frac{1}{2}\right)\left(H_{XY}+\frac{3}{2}\right)}.$$

In particular, $C_Y^\kappa > 0$, while C_{XY}^κ has the sign of ρc and may vanish if $\rho c = 0$.

Proof. Insert (12) into (13) and integrate in time:

$$\int_0^T t^{H+\frac{1}{2}}\, dt = \frac{T^{H+\frac{3}{2}}}{H+\frac{3}{2}}.$$

This gives

$$\kappa_3(T) = 6\sigma_0^3 \left(\frac{a_Y}{H_Y+\frac{3}{2}} T^{H_Y+\frac{3}{2}} + \frac{a_{XY}}{H_{XY}+\frac{3}{2}} T^{H_{XY}+\frac{3}{2}} \right) + o(T^{H_*+\frac{3}{2}}),$$

which is exactly the stated formula. □

7.4. **From the third cumulant to ATM Bachelier skew.**

Lemma 7.6 (Variance expansion of the Bachelier return)**.** *Let*

$$V_T := \mathbb{E}[R_T^2].$$

Then, as $T \downarrow 0$,

$$V_T = \sigma_0^2 T + O\Big(T^{2H_*+1}\Big).$$

Consequently,

$$\frac{\sigma_0\sqrt{T}}{\sqrt{V_T}} = 1 + O(T^{2H_*}).$$

Proof. Set

$$\Sigma_t^2 := \mathrm{Var}(V_t + c\nu_Y Z_t).$$

By Lemma 7.1,

$$\Sigma_t^2 = O(t^{2H_*}) \qquad \text{uniformly for } t \le T.$$

Since $V_t + c\nu_Y Z_t$ is centred Gaussian and $C(t)$ is the convexity correction chosen so that $\mathbb{E}[e^{\mathcal{Y}_t}] = 1$, we have

$$C(t) = \frac{1}{2}\Sigma_t^2.$$

Therefore

$$\mathbb{E}[e^{2\mathcal{Y}_t}] = e^{-2C(t)}\mathbb{E}[e^{2(V_t+c\nu_Y Z_t)}] = e^{\Sigma_t^2} = 1 + O(t^{2H_*}),$$

uniformly for $t \le T$. By Itô isometry,

$$V_T = \sigma_0^2 \int_0^T \mathbb{E}[e^{2\mathcal{Y}_t}]\, dt = \sigma_0^2 \int_0^T (1 + O(t^{2H_*}))\, dt = \sigma_0^2 T + O\Big(T^{2H_*+1}\Big).$$

Finally,

$$\frac{V_T}{\sigma_0^2 T} = 1 + O(T^{2H_*}),$$

and the expansion of $(1+x)^{-1/2}$ near $x = 0$ yields

$$\frac{\sigma_0\sqrt{T}}{\sqrt{V_T}} = 1 + O(T^{2H_*}).$$ □

Define the standardised return

$$\widetilde{R}_T := \frac{R_T}{\sqrt{V_T}}, \qquad \gamma_3(T) := \mathbb{E}[\widetilde{R}_T^3] = \frac{\kappa_3(T)}{V_T^{3/2}}.$$

Let $\psi(T)$ denote the at-the-money Bachelier implied-volatility skew, i.e.

$$\psi(T) := \partial_k \sigma_{\mathrm{N}}(T,k)\big|_{k=0},$$

where $\sigma_{\mathrm{N}}(T,k)$ is the normal implied volatility at maturity T and strike k.

Proposition 7.7 (Small-chaos structure of the standardised return)**.** *Let*

$$Z_T := \frac{B_T}{\sqrt{T}}, \qquad U_T := \frac{1}{\sqrt{T}} \int_0^T \mathcal{Y}_t \, dB_t.$$

Then $Z_T \sim N(0,1)$, and for each fixed $m \geq 1$,

$$\mathbb{E}[|U_T|^m] = O(T^{mH_*}).$$

In addition,

$$\widetilde{R}_T = Z_T + U_T + \widetilde{\mathcal{R}}_T,$$

where

$$\mathbb{E}[|\widetilde{\mathcal{R}}_T|^3] = o(T^{H_*}),$$

and U_T belongs to the Wiener chaoses of order at most two generated by $(W^X, W^\perp)$. In particular,

$$\widetilde{R}_T - Z_T \to 0 \quad \textit{in } L^m$$

for every fixed $m \geq 1$, and

$$\gamma_3(T) = O(T^{H_*}).$$

Proof. By Lemma 7.6,

$$\alpha_T := \frac{\sigma_0 \sqrt{T}}{\sqrt{V_T}} - 1 = O(T^{2H_*}).$$

Dividing the decomposition from Lemma 7.3 by $\sqrt{V_T}$ gives

$$\widetilde{R}_T = \frac{\sigma_0 B_T}{\sqrt{V_T}} + \frac{\sigma_0}{\sqrt{V_T}} \int_0^T \mathcal{Y}_t \, dB_t + \frac{\mathcal{R}_T}{\sqrt{V_T}}.$$

Define

$$\widetilde{\mathcal{R}}_T := \alpha_T Z_T + \alpha_T U_T + \frac{\mathcal{R}_T}{\sqrt{V_T}}.$$

Then

$$\widetilde{R}_T = Z_T + U_T + \widetilde{\mathcal{R}}_T.$$

By Burkholder–Davis–Gundy and Lemma 7.1,

$$\mathbb{E}[|U_T|^m] \leq C_m \, T^{-m/2} \mathbb{E}\left[\left(\int_0^T |\mathcal{Y}_t|^2 \, dt \right)^{m/2} \right] = O(T^{mH_*}).$$

Since $\mathbb{E}[|Z_T|^3] = O(1)$, $\mathbb{E}[|U_T|^3] = O(T^{3H_*})$, and $\alpha_T = O(T^{2H_*})$, we have

$$\mathbb{E}[|\alpha_T Z_T + \alpha_T U_T|^3] = O(T^{6H_*}) = o(T^{H_*}).$$

Likewise, Lemma 7.3 and Lemma 7.6 give

$$\mathbb{E}[|\widetilde{\mathcal{R}}_T|^3] \leq C \, \mathbb{E}[|\alpha_T Z_T + \alpha_T U_T|^3] + C \, V_T^{-3/2} \mathbb{E}[|\mathcal{R}_T|^3] = o(T^{H_*}).$$

The process $\mathcal{Y}_t$ is the sum of a deterministic term and first-chaos Gaussian Volterra functionals, so $\int_0^T \mathcal{Y}_t \, dB_t$ belongs to the Wiener chaoses of order at most two. The L^m-convergence $\widetilde{R}_T - Z_T \to 0$ follows from the bounds on U_T and $\widetilde{\mathcal{R}}_T$. Finally,

$$\gamma_3(T) = \frac{\kappa_3(T)}{V_T^{3/2}} = O(T^{H_*})$$

by Theorem 7.5 and $V_T \sim \sigma_0^2 T$. $\square$

Theorem 7.8 (ATM Bachelier skew–cumulant transfer). *In the present Gaussian–Volterra model,*

$$\psi(T) = \frac{1}{6\sqrt{T}}\,\gamma_3(T) + o\Big(T^{H_*-\frac{1}{2}}\Big)\,.$$

Equivalently,

$$\psi(T) = \frac{1}{6\sigma_0^3 T^2}\,\kappa_3(T) + o\Big(T^{H_*-\frac{1}{2}}\Big)\,.$$

Proof. Let

$$X_T := \widetilde{R}_T.$$

By Proposition 7.7, X_T is a standard Gaussian core plus a finite-chaos perturbation that is L^m-small of order T^{H_*}, with $\gamma_3(T) \to 0$. The non-degenerate Gaussian core gives a smooth local density, and the perturbation remains in a fixed finite sum of Wiener chaoses at the required order. This is the one-term Edgeworth regime for smooth perturbations of a Gaussian law; see Bhattacharya and Rao [28], Chapter 2, and the short-time rough-volatility discussion in Fukasawa [25], §3. The distribution function of X_T therefore satisfies

$$\mathbb{Q}(X_T \le x) = \Phi(x) - \frac{\gamma_3(T)}{6} H_2(x)\phi(x) + o(T^{H_*}),$$

uniformly for x in bounded sets, where ϕ and Φ denote the standard normal density and distribution function, and $H_2(x) = x^2 - 1$.

Evaluating at $x = 0$ gives

$$\mathbb{Q}(X_T > 0) = \frac{1}{2} - \frac{\gamma_3(T)}{6\sqrt{2\pi}} + o(T^{H_*}),$$

because $H_2(0) = -1$ and $\phi(0) = 1/\sqrt{2\pi}$.

Now

$$\partial_k C(T,k)\big|_{k=0} = -\mathbb{Q}(R_T > 0) = -\mathbb{Q}(X_T > 0) = -\frac{1}{2} + \frac{\gamma_3(T)}{6\sqrt{2\pi}} + o(T^{H_*}),$$

for the market call price $C(T,k) := \mathbb{E}[(R_T - k)^+]$.

For the Bachelier model with volatility parameter σ, the ATM call price derivative is

$$\partial_k C_{\mathrm{N}}(T,k;\sigma)\big|_{k=0} = -\frac{1}{2},$$

independently of σ, and the ATM vega is

$$\partial_\sigma C_{\mathrm{N}}(T,0;\sigma) = \sqrt{T}\,\phi(0) = \frac{\sqrt{T}}{\sqrt{2\pi}}.$$

Implicit differentiation of

$$C(T,k) = C_{\mathrm{N}}(T,k;\sigma_{\mathrm{N}}(T,k))$$

at $k = 0$, evaluated at the ATM value $\sigma_{\mathrm{N}}(T,0)$, therefore yields

$$\psi(T) = \frac{\partial_k C(T,0) - \partial_k C_{\mathrm{N}}(T,0;\sigma_{\mathrm{N}}(T,0))}{\partial_\sigma C_{\mathrm{N}}(T,0;\sigma_{\mathrm{N}}(T,0))} = \frac{\gamma_3(T)}{6\sqrt{T}} + o\Big(T^{H_*-\frac{1}{2}}\Big)\,.$$

Since $\gamma_3(T) = \kappa_3(T)/V_T^{3/2}$, Lemma 7.6 yields

$$\frac{1}{6\sqrt{T}}\gamma_3(T) = \frac{1}{6\sigma_0^3 T^2}\kappa_3(T) + o\Big(T^{H_*-\frac{1}{2}}\Big)\,,$$

which gives the equivalent cumulant form. □

Remark 7.9 (Why the pricing transfer is model-native). The pricing block uses only one external ingredient, the standard one-term Edgeworth transfer from a small perturbation of a Gaussian core to the digital price at the money. All other terms are derived inside the model. Proposition 7.7 verifies the required perturbative form, with an explicit first-chaos Gaussian core, an at-most-second-chaos correction of size T^{H_*}, and a smaller remainder. The $\mathbb{Q}$-threshold is therefore the pricing projection of the same Gaussian-Volterra short-time structure that drives the cumulant expansion.

Remark 7.10 (Relative remainder under non-cancellation). If the leading coefficient in Theorem 7.5 does not cancel, so that

$$\kappa_3(T) \asymp T^{H_*+\frac{3}{2}} \qquad \text{and hence} \qquad \gamma_3(T) \asymp T^{H_*},$$

then the absolute remainder in Theorem 7.8 can be rewritten in the relative form

$$\psi(T) = \frac{1}{6\sqrt{T}}\,\gamma_3(T) + o\!\left(\frac{\gamma_3(T)}{\sqrt{T}}\right).$$

The absolute version is the one that remains valid uniformly, including the exceptional cancellation regime.

Theorem 7.11 (Sharp pricing-level visibility threshold)**.** *As $T \downarrow 0$,*

$$\psi(T) = \widetilde{C}_Y\,T^{H_Y-\frac{1}{2}} + \widetilde{C}_{XY}\,T^{H_{XY}-\frac{1}{2}} + o\Big(T^{H_*-\frac{1}{2}}\Big),$$

where

$$\widetilde{C}_Y := \frac{C_Y^\kappa}{6\sigma_0^3} = \frac{\nu_Y\sqrt{1-\rho^2}\,L_Y(0,0)}{\left(H_Y+\frac{1}{2}\right)\left(H_Y+\frac{3}{2}\right)},$$

and

$$\widetilde{C}_{XY} := \frac{C_{XY}^\kappa}{6\sigma_0^3} = \frac{c\nu_Y\rho\,L_{XY}(0,0)}{\left(H_{XY}+\frac{1}{2}\right)\left(H_{XY}+\frac{3}{2}\right)}.$$

Note that $\widetilde{C}_Y > 0$, while

$$\widetilde{C}_{XY} \neq 0 \quad \Longleftrightarrow \quad \rho\,c\,L_{XY}(0,0) \neq 0.$$

Accordingly:

(1) If $H_{XY} > H_Y$, the cross-asset component is asymptotically negligible at leading order, and

$$\psi(T) \sim \widetilde{C}_Y\,T^{H_Y-\frac{1}{2}}.$$

(2) If $H_{XY} < H_Y$ and $\widetilde{C}_{XY} \neq 0$, the cross-asset component determines the leading-order skew, and

$$\psi(T) \sim \widetilde{C}_{XY}\,T^{H_{XY}-\frac{1}{2}}.$$

(3) If $H_{XY} = H_Y$, both components enter at the same order:

$$\psi(T) = (\widetilde{C}_Y + \widetilde{C}_{XY})\,T^{H_Y-\frac{1}{2}} + o\Big(T^{H_Y-\frac{1}{2}}\Big).$$

Hence, outside the exceptional cancellation case $H_{XY} = H_Y$ and $\widetilde{C}_Y + \widetilde{C}_{XY} = 0$, the sharp pricing-level threshold is

$$H_{XY} = H_Y.$$

Equivalently, provided $\widetilde{C}_{XY} \neq 0$, the cross-asset component contributes at leading order to the ATM Bachelier skew if and only if

$$H_{XY} \leq H_Y.$$

In the boundary case $H_{XY} = H_Y$, this leading-order statement is understood outside the exceptional cancellation regime $\widetilde{C}_Y + \widetilde{C}_{XY} = 0$.

Proof. Combine Theorems 7.5 and 7.8. Since

$$\psi(T) = \frac{1}{6\sigma_0^3 T^2}\,\kappa_3(T) + o\Big(T^{H_*-\frac{1}{2}}\Big),$$

the powers of T in the skew expansion are obtained by subtracting 2 from the powers in the cumulant expansion. The three cases then follow by comparing

$$H_Y - \frac{1}{2} \qquad \text{and} \qquad H_{XY} - \frac{1}{2}.$$

□

Theorem 7.11 separates pricing visibility from the presence of the cross-asset kernel. The term with the smaller Hurst exponent controls the dominant short-maturity skew. A smoother cross component remains outside the leading pricing projection unless the exponents coincide.

8. $\mathbb{P}$-Measure: Path Detectability

Corollary 8.1 (Path detectability in the smoothing regime)**.** *Under the assumptions of Theorem 6.11 with $H_{XY} > H_Y$, the coupled and uncoupled centred Gaussian path laws are*

(i) mutually singular if

$$H_Y < H_{XY} \leq H_Y + \tfrac{1}{4};$$

(ii) equivalent if

$$H_{XY} > H_Y + \tfrac{1}{4}.$$

Proof. This is the path-space formulation of Theorem 6.11 in the smoothing regime $H_{XY} > H_Y$. □

Remark 8.2 (Relative-entropy interpretation). In the equivalent regime $H_{XY} > H_Y + \frac{1}{4}$, one may define the path-space Kullback–Leibler divergence between the coupled and uncoupled centred Gaussian measures. Its finiteness is controlled by the same relative covariance operator $\mathcal{R}_{Y,XY}$ from Theorem 6.8. In the singular regime, the path-space Kullback–Leibler divergence is infinite.

Remark 8.3 (Discrete sampling caution). A discretised information criterion based directly on fine-scale fractional increments need not reflect the path-space divergence. Rough paths with $H < \frac{1}{2}$ have singular small-scale behaviour, so a finite-resolution proxy is a heuristic surrogate rather than a substitute for the path-space comparison.

Remark 8.4 (Finite data). Finite-resolution observation replaces the path-space dichotomy by a sequence of statistical experiments. Path-space singularity in the latent regime does not produce an infinite statistic in any finite dataset. It means that increasingly fine observation can separate the coupled and uncoupled models with increasing power.

9. The latent contagion regime

Proposition 9.1 (Threshold summary)**.** *Assume the skew expansion in Section 7 with nonzero leading coefficients, and assume $c > 0$. The pricing statement is read from Theorem 7.11. The path-law statements are asserted only in the smoothing regime $H_{XY} > H_Y$, where Theorem 6.11 applies. Then the pricing-level and path-level thresholds are*

$$H_{XY}^{\mathbb{Q}} = H_Y, \qquad H_{XY}^{\mathbb{P}} = H_Y + \tfrac{1}{4}.$$

Accordingly:

Regime	*Leading-order skew visibility*	*Path-law relation*
$H_{XY} \leq H_Y$	*visible*	*(not established here)*
$H_Y < H_{XY} \leq H_Y + \frac{1}{4}$	*not visible at leading order*	*singular*
$H_{XY} > H_Y + \frac{1}{4}$	*not visible at leading order*	*equivalent*

Proof. The pricing threshold follows by comparing the powers in the skew expansion of Theorem 7.11. Once $H_{XY} > H_Y$, the cross term is smoother than the idiosyncratic term and drops out of the leading skew. In that same smoothing regime, Corollary 8.1 gives singularity for $H_Y < H_{XY} \leq H_Y + \frac{1}{4}$ and equivalence for $H_{XY} > H_Y + \frac{1}{4}$. The row $H_{XY} \leq H_Y$ is therefore a pricing classification with no modulated path-law assertion here. □

Remark 9.2 (Status of the non-smoothing regime). The classification distinguishes three regimes by relative Hurst exponent. The path-law statements are proved in the smoothing regime $H_{XY} > H_Y$. For $H_{XY} \leq H_Y$, the pure fractional setting is consistent with mutual singularity, but the corresponding result for the present modulated model is not established here.

Remark 9.3 (Existence). Under the standing range $H_{XY} < \frac{1}{2}$, the latent interval is

$$(H_Y,\ \min\{H_Y + \tfrac{1}{4}, \tfrac{1}{2}\}).$$

It is nonempty for every $H_Y \in (0, \frac{1}{2})$. If $H_Y < \frac{1}{4}$, the full quarter-gap interval is available. For empirically rough values such as $H_Y \approx 0.1$ [1], the admissible cross-exponent interval is nontrivial.

Remark 9.4 (Origin of the gap). The 1/4 gap is controlled by the relative factor $B = C_Y^{-1/2} C_{XY}^{1/2}$. After factorisation through the rougher kernel, B has the same Schatten-4 membership as the fractional integral $I_{0+}^{H_{XY}-H_Y}$. Since $s_n(I_{0+}^{\delta}) \asymp n^{-\delta}$, the summability criterion becomes $\sum_n n^{-4(H_{XY}-H_Y)} < \infty$, which is equivalent to $H_{XY} - H_Y > \frac{1}{4}$. Because $\mathcal{R}_{Y,XY} = BB^*$, this is exactly the Hilbert–Schmidt threshold for $\mathcal{R}_{Y,XY}$.

Remark 9.5 (Interpretation). In the latent regime, the contagion kernel $(t-s)^{H_{XY}-1/2} L_{XY}(t,s)$ is smoother than the idiosyncratic kernel $(t-s)^{H_Y-1/2} L_Y(t,s)$. It therefore does not dominate the leading short-maturity skew. The same kernel is not smooth enough to be absorbed at the path-law level. Pricing visibility stops at the leading skew projection; path-law distinguishability still records the residual Volterra perturbation.

9.1. **Time-varying amplitudes.** If the contagion amplitude $\eta_{X\to Y}$ varies over time, the skew crossover scale changes. The threshold classification derived here depends on the exponents H_Y and H_{XY}, provided the contagion amplitude does not vanish identically.

10. Spectral structure and latent contagion

Section 6 characterises the covariance perturbation

$$\mathcal{R}_{Y,XY} = C_Y^{-1/2} C_{XY} C_Y^{-1/2}$$

through the Feldman–Hájek criterion in the smoothing regime $H_{XY} > H_Y$. The resulting path-space detectability threshold is

$$H_{XY} = H_Y + \tfrac{1}{4}.$$

The spectral asymptotics in Theorem 6.8 refine that binary classification.

Spectral asymptotics. Under the assumptions of Theorem 6.8, the eigenvalues of $\mathcal{R}_{Y,XY}$ satisfy

$$\lambda_n(\mathcal{R}_{Y,XY}) \asymp n^{-2\delta}, \qquad \delta = H_{XY} - H_Y > 0.$$

The perturbation has a power-law spectral profile determined by the gap δ. Hilbert–Schmidt summability is one boundary in that profile, not the disappearance of latent spectral mass.

Definition 10.1 (Latent spectral energy)**.** For $N \geq 1$, define the cumulative spectral mass

$$\mathcal{E}_{\text{latent}}(N) := \sum_{n=1}^{N} \lambda_n(\mathcal{R}_{Y,XY}).$$

By Corollary 6.13, this quantity satisfies

$$\mathcal{E}_{\text{latent}}(N) \asymp \begin{cases} N^{1-2\delta}, & 0 < \delta < \frac{1}{2}, \\ \log N, & \delta = \frac{1}{2}, \\ 1, & \delta > \frac{1}{2}. \end{cases}$$

Relation to detectability. The spectral asymptotics and the equivalence criterion separate three cases.

- For $0 < \delta \leq \frac{1}{4}$, the operator $\mathcal{R}_{Y,XY}$ is not Hilbert–Schmidt and the latent spectral mass diverges.
- For $\frac{1}{4} < \delta < \frac{1}{2}$, the path laws are equivalent, but the cumulative spectral mass still diverges.
- For $\delta > \frac{1}{2}$, the latent spectral energy remains bounded.

The detectability threshold $\delta = \frac{1}{4}$ is not the boundary at which cumulative spectral mass disappears.

Interpretation. After normalisation by $C_Y^{1/2}$, the cross-asset covariance contribution is represented by a hierarchy of orthogonal directions whose strengths decay according to a fractional power law. The operator

$$T = C_Y^{-1/2} C_{XY} C_Y^{-1/2}$$

therefore encodes the latent spectral signature of the cross-asset covariance contribution in the volatility field of asset Y.

Resolution dependence. At temporal resolution Δ, the effective number of resolved modes satisfies $N_\Delta \asymp \Delta^{-1}$. Consequently,

$$\mathcal{E}_{\text{latent}}(N_\Delta) \asymp \Delta^{-(1-2\delta)}, \qquad 0 < \delta < \tfrac{1}{2},$$

so finer observation scales resolve an increasing accumulation of cross-asset spectral mass whenever $\delta < \frac{1}{2}$, even in regimes where the leading short-maturity pricing signal is absent.

Dynamic interpretation. The focus here is detectability and spectral structure at the level of Gaussian path measures. The latent spectral representation also has a dynamic interpretation in terms of rough latent mechanisms. The threshold identified here separates static path-space detectability from the observable-screening and feasible-inference questions required in applications.

Section 11 records finite-dimensional diagnostics for these spectral properties.

11. Finite-dimensional diagnostics for pricing and path-space thresholds

The numerical experiments are diagnostics for the two thresholds derived above. The pricing-level threshold is $H_{XY} = H_Y$ from Theorem 7.11. In the smoothing regime $H_{XY} > H_Y$, the path-space threshold is $H_{XY} = H_Y + \frac{1}{4}$ from Theorem 6.11. The experiments do not enter the proofs. They display the scaling laws in a finite-dimensional Riemann–Liouville benchmark.

The numerical benchmark uses the pure Riemann–Liouville Gaussian core

$$\mathcal{Y}_t = \nu_Y \int_0^t (t-s)^{H_Y - \frac{1}{2}}\, dW_s^\perp + c\nu_Y \int_0^t (t-s)^{H_{XY} - \frac{1}{2}}\, dW_s^X,$$

where

$$B_t = \rho W_t^X + \sqrt{1-\rho^2}\, W_t^\perp$$

is the price Brownian motion. The parameters are

$$H_Y = 0.10, \qquad \rho = -0.70, \qquad c\nu_Y = 0.35, \qquad T = 1.$$

The two principal thresholds are then

$$H_{XY}^{\mathbb{Q}} = H_Y = 0.10, \qquad H_{XY}^{\mathbb{P}} = H_Y + \frac{1}{4} = 0.35.$$

11.1. Pricing-side threshold under $\mathbb{Q}$.

By Theorem 7.5, the third cumulant of the Bachelier return satisfies

$$\kappa_3(T) = C_Y^\kappa\, T^{H_Y + \frac{3}{2}} + C_{XY}^\kappa\, T^{H_{XY} + \frac{3}{2}} + o\Big(T^{H_* + \frac{3}{2}}\Big),$$

and the ATM Bachelier skew scales as

$$\psi(T) \sim T^{\min(H_Y,\, H_{XY}) - \frac{1}{2}}.$$

For pure Riemann–Liouville kernels, the leverage covariance $\mathbb{E}[B_t \mathcal{Y}_t]$, the third cumulant, and the resulting leading-order Bachelier skew are available in closed form (Lemma 7.2 and Theorem 7.5). The diagnostic is therefore deterministic; no Monte Carlo simulation is used.

Figure 1 displays $|\psi(T)|$ on logarithmic axes for four values of H_{XY}. When $H_{XY} < H_Y$, the cross-asset term produces the steeper short-maturity skew. At $H_{XY} = H_Y$, both terms contribute at the same order. When $H_{XY} > H_Y$, the idiosyncratic benchmark governs the asymptotic skew, and the cross-asset component is not visible at leading order to the Bachelier skew.

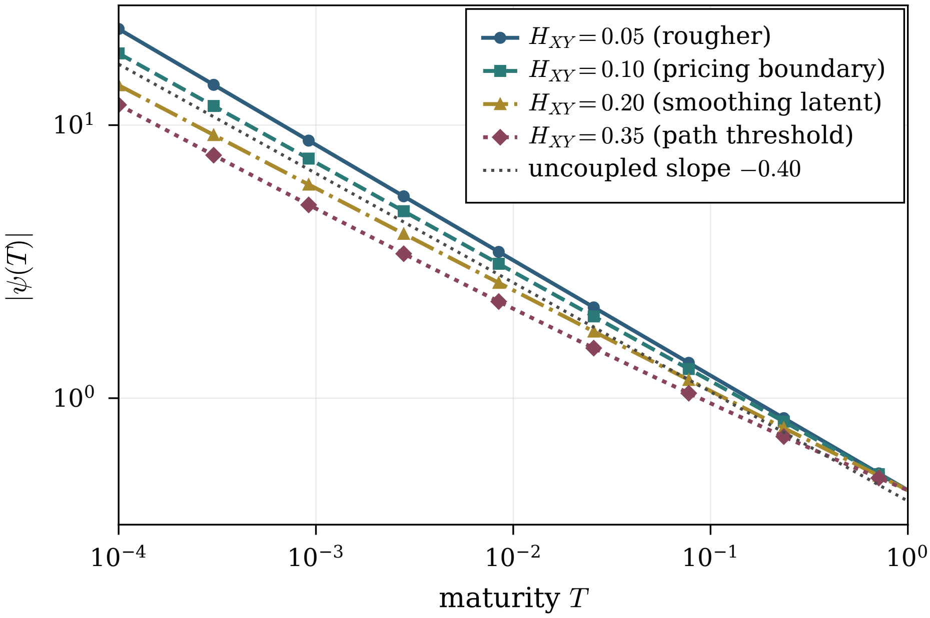

FIGURE 1. Bachelier ATM skew scaling $|\psi(T)|$ for several values of H_{XY}. The dominant short-maturity exponent is $\min(H_Y, H_{XY}) - \frac{1}{2}$. The dashed line is the uncoupled benchmark with slope $H_Y - \frac{1}{2} = -0.40$.

11.2. **Path-space threshold under $\mathbb{P}$.** The path-space calculation uses the discrete analogue of the relative covariance operator,

$$\mathcal{R}_N = C_{Y,N}^{-1/2}\, C_{XY,N}\, C_{Y,N}^{-1/2},$$

constructed from the covariance matrices of the discretised Volterra kernels on a grid of size N. The continuum theory identifies the Hilbert–Schmidt condition

$$\mathcal{R}_{Y,XY} \in \mathfrak{S}_2 \qquad \Longleftrightarrow \qquad H_{XY} > H_Y + \frac{1}{4}.$$

In finite dimension, the corresponding diagnostic is $\mathrm{tr}(\mathcal{R}_N^2)$.

Figure 2 shows $\mathrm{tr}(\mathcal{R}_N^2)$ as the grid size increases, for the four representative values

$$H_{XY} \in \{0.10, 0.25, 0.35, 0.45\}.$$

For $H_{XY} = 0.10$ and $H_{XY} = 0.25$, the quantity grows with N, consistent with the non-Hilbert–Schmidt regime. For $H_{XY} = 0.35$, the critical threshold, the behaviour is borderline over the resolved range. For $H_{XY} = 0.45$, the supercritical case, the growth is weaker. This is the finite-dimensional trace of the path-space threshold in the smoothing regime.

The finite-dimensional likelihood classifier compares the coupled and uncoupled Gaussian laws (Figure 3). Discrimination is strongest in the non-smoothing benchmark $H_{XY} = H_Y$. Near and above the $\mathbb{P}$-threshold, the finite-grid accuracy deteriorates over the grid sizes considered. This is consistent with the Feldman–Hájek dichotomy, since equivalence of path laws precludes perfect separation on the limiting path experiment.

11.3. **Spectral asymptotics and latent energy.** The theory predicts the eigenvalue asymptotics

$$\lambda_n(\mathcal{R}_{Y,XY}) \asymp n^{-2\delta}, \qquad \delta = H_{XY} - H_Y > 0,$$

and cumulative spectral mass

$$\mathcal{E}_{\mathrm{latent}}(N) \asymp N^{1-2\delta}.$$

Figures 4 and 5 display the ordered eigenvalues of the finite relative covariance matrix on logarithmic axes for $H_{XY} = 0.25$ and $H_{XY} = 0.35$, together with the theoretical reference slope. In both cases the numerical spectra follow the predicted power-law decay over the resolved range of modes.

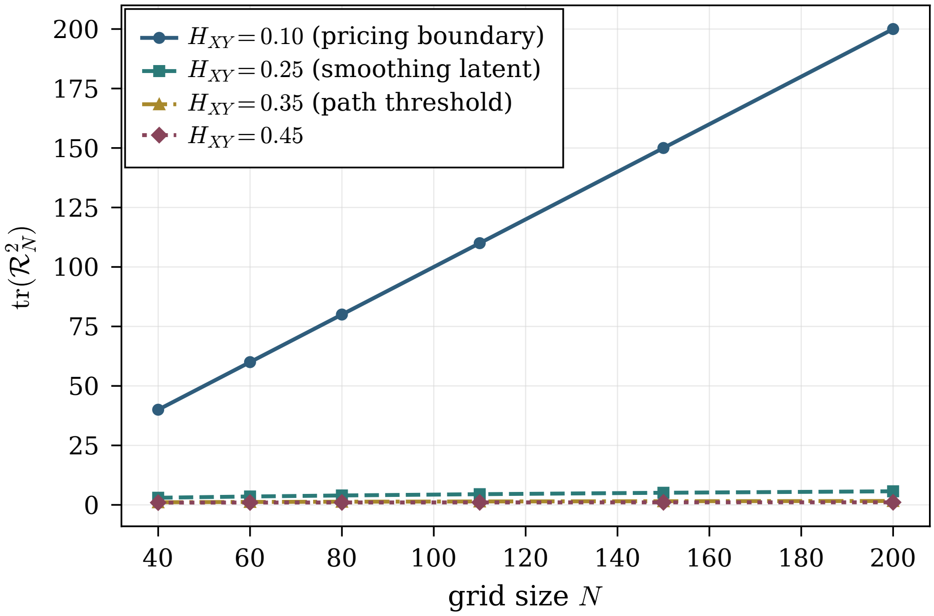

FIGURE 2. Finite-dimensional proxy for the $\mathbb{P}$-threshold. The quantity $\mathrm{tr}(\mathcal{R}_N^2)$ grows with N in the non-Hilbert–Schmidt regime $\delta < \frac{1}{4}$, exhibits borderline behaviour at the critical threshold $\delta = \frac{1}{4}$, and shows weaker growth in the supercritical regime $\delta > \frac{1}{4}$.

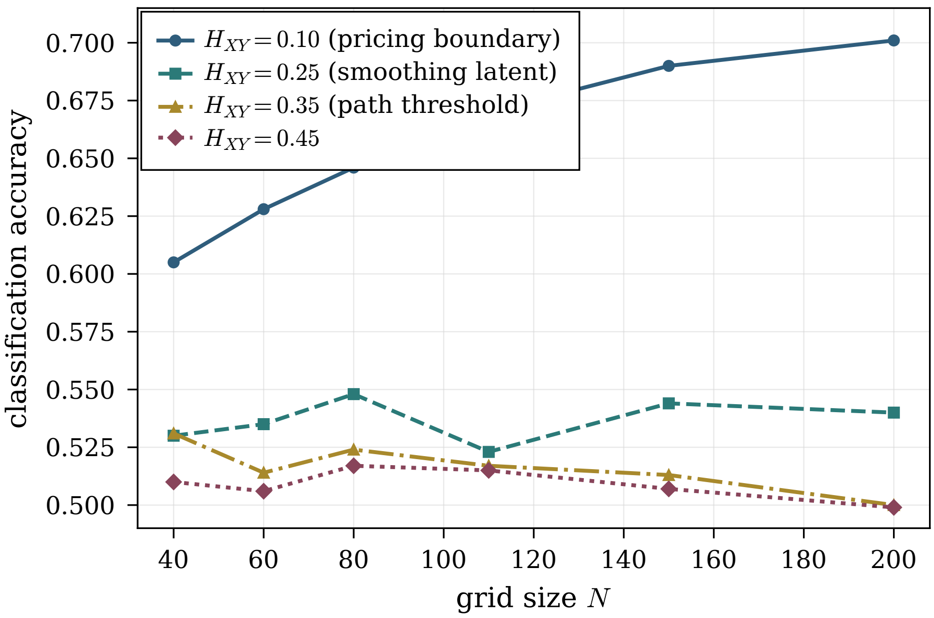

FIGURE 3. Likelihood-based classification accuracy for coupled versus uncoupled Gaussian path laws.

Figures 6 and 7 show that the cumulative latent spectral mass follows the predicted power law over the resolved range. For $H_{XY} = 0.25$, the growth is approximately $N^{0.7}$. For $H_{XY} = 0.35$, the critical threshold, the growth is approximately $N^{1/2}$.

11.4. **Summary.** The numerical diagnostics record the threshold separation:

(1) On the *pricing side*, the Bachelier ATM skew is dominated by the term with the smaller Hurst exponent. Once $H_{XY} > H_Y$, the cross-asset contribution is asymptotically invisible

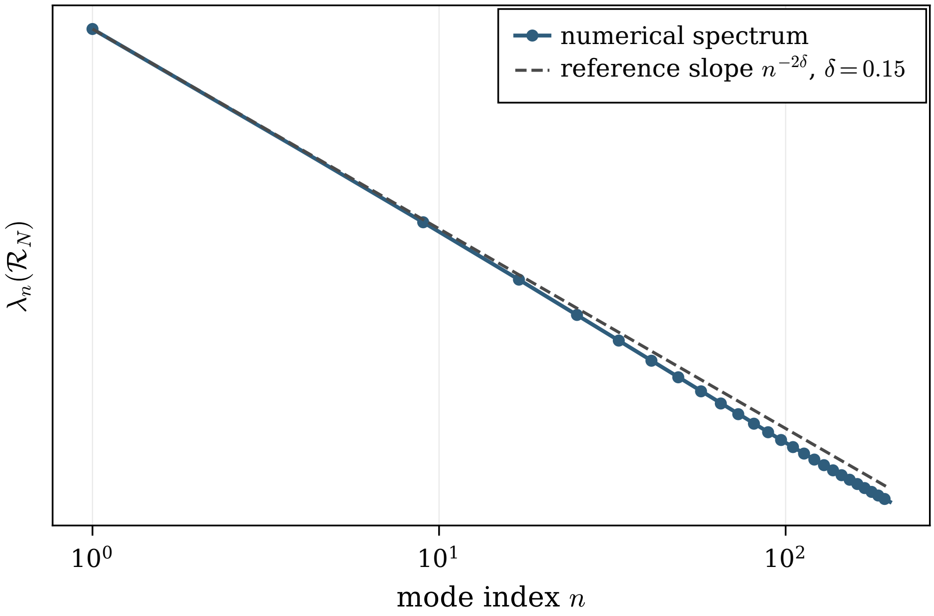

FIGURE 4. Spectrum of $\mathcal{R}_N$ for $H_{XY} = 0.25$ ($\delta = 0.15$). The eigenvalues follow the predicted decay $\lambda_n(\mathcal{R}_{Y,XY}) \asymp n^{-2\delta}$.

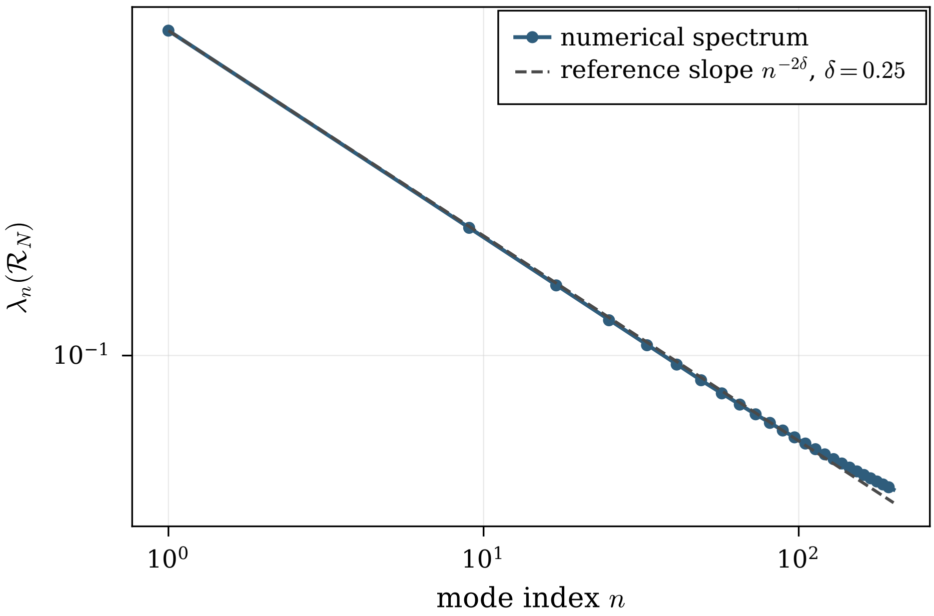

FIGURE 5. Spectrum of $\mathcal{R}_N$ for $H_{XY} = 0.35$ ($\delta = 0.25$), corresponding to the critical $\mathbb{P}$-threshold.

at leading order. The threshold is

$$H_{XY}^{\mathbb{Q}} = H_Y.$$

(2) On the *path-space side*, the relative covariance operator $\mathcal{R}_{Y,XY}$ fails to be Hilbert–Schmidt when $\delta \leq \frac{1}{4}$, and the cumulative spectral mass grows according to the predicted power law. The Feldman–Hájek threshold is

$$H_{XY}^{\mathbb{P}} = H_Y + \frac{1}{4}.$$

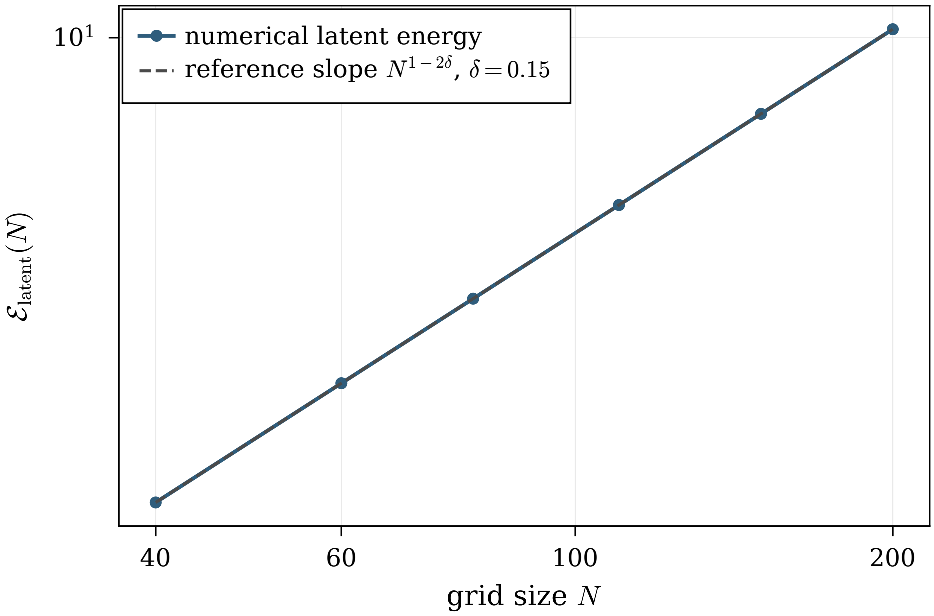

FIGURE 6. Latent spectral energy for $H_{XY} = 0.25$ ($\delta = 0.15$). The growth agrees with the predicted scaling $N^{1-2\delta}$.

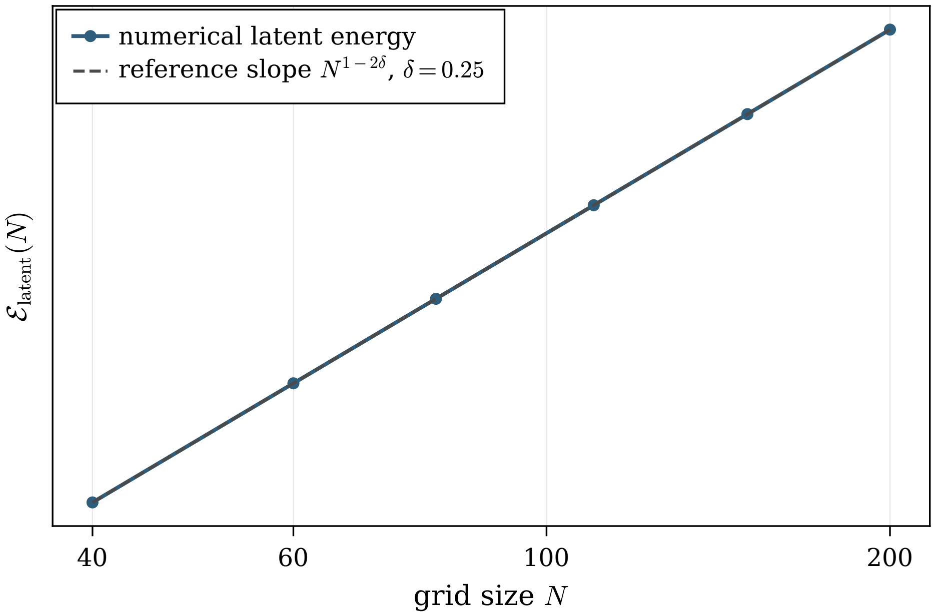

FIGURE 7. Latent spectral energy for $H_{XY} = 0.35$ ($\delta = 0.25$), corresponding to the critical $\mathbb{P}$-threshold. The energy grows as $N^{1/2}$.

(3) In the *latent regime*

$$H_Y < H_{XY} \leq H_Y + \frac{1}{4},$$

the pricing-side diagnostics show an invisible cross-asset component, while the path-space diagnostics exhibit persistent spectral mass and nontrivial growth of the discrete relative-covariance proxies. The finite benchmark therefore separates pricing information from statistical path information.

12. Discussion

12.1. **Relation to existing multivariate rough volatility models.** The model (1)–(3) is a bivariate Volterra stochastic volatility model with heterogeneous Hurst exponents. Dugo, Giorgio, and Pigato [7] study a multivariate fractional Ornstein–Uhlenbeck model with component-wise Hurst exponents and estimate asymmetric cross-covariance spillovers by GMM on realised-volatility data. The classification here has a different object: Gaussian path-law distinguishability rather than spillover magnitude.

Bibinger, Yu, and Zhang [9] find empirically that multivariate fractional Brownian motion forecasts outperform univariate ones when Hurst exponents diverge across assets. That observation is consistent with the threshold structure. When $H_{XY} \neq H_Y$, the cross-asset component may carry information at a roughness scale different from the idiosyncratic one.

12.2. **Relation to the correlation risk premium.** The divergence between $\mathbb{Q}$-measure and $\mathbb{P}$-measure contagion thresholds is related to a separate empirical boundary between implied and realised dependence. Driessen, Maenhout, and Vilkov [29] show that option-implied correlation systematically exceeds realised correlation, with the difference attributed to a large negative correlation risk premium. Chen, Wang *et al.* [30] construct multiplex spillover networks from implied volatility, realised volatility, and variance-risk-premium channels, and find non-redundant contagion information across timescales.

The threshold comparison gives one formal route through which such $\mathbb{Q}$–$\mathbb{P}$ divergence can arise. If the cross-Hurst exponent lies in the latent regime, the $\mathbb{Q}$-measure surface can be insensitive at leading order to the contagion while $\mathbb{P}$-measure path statistics are not. This is not a claim that the empirical correlation risk premium is caused by the latent regime. The claim is narrower: threshold separation is one channel through which $\mathbb{Q}$- and $\mathbb{P}$-measure contagion assessments can disagree. Latent contagion then reflects a separation between statistical and pricing observability rather than an additional risk factor.

12.3. **Limitations and caveats.** The model assumes log-normal volatility, Gaussian driving noise, and a time-invariant cross-exponent H_{XY}. Empirical Hurst exponents may vary over time and may themselves be estimated with error. Cont and Das [31] show that microstructure noise can make standard volatility appear rougher than it is, which complicates estimation of both H_Y and H_{XY}. The threshold picture is a model statement; empirical use depends on exponent estimation.

The factor $\frac{1}{4}$ in the path-detectability threshold is exact within the model. It should not be read with spurious precision in applications. In finite samples, the boundary between the latent and absorbed regimes is not sharp. Detectability depends on sample size, sampling frequency, and noise level. The analysis also relies on Gaussian Volterra structure; non-Gaussian or jump-driven extensions require different path-law tools.

13. Conclusion

The bivariate rough Volterra experiment separates two observation levels for a directed cross-asset volatility channel. The pricing threshold

$$H_{XY}^{\mathbb{Q}} = H_Y$$

governs leading-order Bachelier skew visibility. The path-space threshold

$$H_{XY}^{\mathbb{P}} = H_Y + \tfrac{1}{4}$$

governs equivalence versus singularity of the coupled and uncoupled Gaussian laws in the smoothing regime. The interval between them is the latent contagion regime: the directed channel is absent from leading-order prices and still present in the physical path law.

Smooth non-degenerate modulation does not remove this separation. Theorem 6.8 shows that, for $c > 0$, the relative covariance operator has spectral order $\lambda_n(\mathcal{R}_{Y,XY}) \asymp n^{-2\delta}$ with $\delta = H_{XY} - H_Y > 0$. The fractional singular-value order that governs the pure benchmark therefore continues to govern the modulated Volterra model class used in practice.

The theorem remains a static threshold statement. Observable screening, sequential revelation, and feasible inference belong to separate experiments. The Feldman–Hájek comparison proved here is local to the smoothing regime $H_{XY} > H_Y$, and finite samples replace the path-space dichotomy by a sequence of statistical experiments.

Two residual questions remain. The first is how far the smooth-modulation hypotheses can be weakened without changing the threshold order. The second is whether an analogue survives outside Gaussian Volterra structure, where the Feldman–Hájek mechanism no longer decides path-law comparison. Pricing visibility and path-law detectability are different projections of the same rough channel.

References

[1] J. Gatheral, T. Jaisson, and M. Rosenbaum, Volatility is rough, *Quantitative Finance* **18**(6):933–949, 2018, doi:10.1080/14697688.2017.1393551.

[2] C. Bayer, P. Friz, and J. Gatheral, Pricing under rough volatility, *Quantitative Finance* **16**(6):887–904, 2016, doi:10.1080/14697688.2015.1099717.

[3] O. El Euch and M. Rosenbaum, The characteristic function of rough Heston models, *Mathematical Finance* **29**(1):3–38, 2019, doi:10.1111/mafi.12173.

[4] E. Abi Jaber, M. Larsson, and S. Pulido, Affine Volterra processes, *The Annals of Applied Probability* **29**(5):3155–3200, 2019, doi:10.1214/19-AAP1477.

[5] G. Livieri, S. Mouti, A. Pallavicini, and M. Rosenbaum, Rough volatility: Evidence from option prices, *IISE Transactions* **50**(9):767–776, 2018, doi:10.1080/24725854.2018.1444297.

[6] M. Tomas and M. Rosenbaum, From microscopic price dynamics to multidimensional rough volatility models, *Advances in Applied Probability* **53**(2):425–462, 2021, doi:10.1017/apr.2020.60.

[7] R. Dugo, G. Giorgio, and P. Pigato, *The multivariate fractional Ornstein–Uhlenbeck process*, Stochastic Processes and their Applications **192** (2026), article 104814, doi:10.1016/j.spa.2025.104814.

[8] F. Lavancier, A. Philippe, and D. Surgailis, Covariance function of vector self-similar processes, *Statistics & Probability Letters* **79**(23):2415–2421, 2009, doi:10.1016/j.spl.2009.08.015.

[9] M. Bibinger, J. Yu, and C. Zhang, Modeling and forecasting realized volatility with multivariate fractional Brownian motion, preprint, arXiv:2504.15985, 2025.

[10] K. J. Forbes and R. Rigobon, No contagion, only interdependence: Measuring stock market comovements, *The Journal of Finance* **57**(5):2223–2261, 2002, doi:10.1111/0022-1082.00494.

[11] M. Pericoli and M. Sbracia, A primer on financial contagion, *Journal of Economic Surveys* **17**(4):571–608, 2003, doi:10.1111/1467-6419.00205.

[12] F. X. Diebold and K. Yılmaz, Better to give than to receive: Predictive directional measurement of volatility spillovers, *International Journal of Forecasting* **28**(1):57–66, 2012, doi:10.1016/j.ijforecast.2011.02.006.

[13] J. Baruník and T. Křehlík, Measuring the frequency dynamics of financial connectedness and systemic risk, *Journal of Financial Econometrics* **16**(2):271–296, 2018, doi:10.1093/jjfinec/nby001.

[14] R. H. Cameron and W. T. Martin, Transformations of Wiener integrals under translations, *Annals of Mathematics* **45**(2):386–396, 1944, doi:10.2307/1969276.

[15] J. Feldman, Equivalence and perpendicularity of Gaussian processes, *Pacific Journal of Mathematics* **8**(4):699–708, 1958, doi:10.2140/pjm.1958.8.699.

[16] J. Hájek, On a property of normal distributions of any stochastic process, *Czechoslovak Mathematical Journal* **8**(4):610–618, 1958, doi:10.21136/CMJ.1958.100333.

[17] V. I. Bogachev, *Gaussian Measures*, Mathematical Surveys and Monographs, Vol. 62, American Mathematical Society, Providence, RI, 1998.

[18] F. Baudoin and D. Nualart, Equivalence of Volterra processes, *Stochastic Processes and Applications* **107**(2):327–350, 2003, doi:10.1016/S0304-4149(03)00088-7.

[19] P. Cheridito, Mixed fractional Brownian motion, *Bernoulli* **7**(6):913–934, 2001, doi:10.2307/3318626.

[20] B. L. S. Prakasa Rao, Singularity of fractional Brownian motions with different Hurst indices, *Stochastic Analysis and Applications* **26**(2):334–337, 2008, doi:10.1080/07362990701857277.

[21] M. Sh. Birman and M. Z. Solomyak, Estimates of singular numbers of integral operators, *Russian Mathematical Surveys* **32**(1):15–89, 1977, doi:10.1070/RM1977v032n01ABEH001592.

[22] M. R. Dostanić, Asymptotic behavior of the singular values of fractional integral operators, *Journal of Mathematical Analysis and Applications* **175**(2):380–391, 1993, doi:10.1006/jmaa.1993.1177.

[23] V. K. Tuan and R. Gorenflo, Asymptotics of singular values of fractional integral operators, *Inverse Problems* **10**(4):949–955, 1994, doi:10.1088/0266-5611/10/4/013.

[24] R. Gorenflo and V. K. Tuan, Singular value decomposition of fractional integration operators in L_2-spaces with weights, *Journal of Inverse and Ill-Posed Problems* **3**(1):1–10, 1995, doi:10.1515/jiip.1995.3.1.1.

[25] M. Fukasawa, Short-time at-the-money skew and rough fractional volatility, *Quantitative Finance* **17**(2):189–198, 2017, doi:10.1080/14697688.2016.1197410.

[26] S. G. Samko, A. A. Kilbas, and O. I. Marichev, *Fractional Integrals and Derivatives: Theory and Applications*, Gordon and Breach, Amsterdam, 1993.

[27] B. Simon, *Trace Ideals and Their Applications*, 2nd ed., Mathematical Surveys and Monographs, Vol. 120, American Mathematical Society, Providence, RI, 2005, doi:10.1090/surv/120.

[28] R. N. Bhattacharya and R. Ranga Rao, *Normal Approximation and Asymptotic Expansions*, Wiley, New York, 1976.

[29] J. Driessen, P. J. Maenhout, and G. Vilkov, The price of correlation risk: Evidence from equity options, *The Journal of Finance* **64**(3):1377–1406, 2009, doi:10.1111/j.1540-6261.2009.01467.x.

[30] J. Gong, G.-J. Wang, Y. Zhou, Y. Zhu, C. Xie, and M. Foglia, Spreading of cross-market volatility information: Evidence from multiplex network analysis of volatility spillovers, *Journal of International Financial Markets, Institutions and Money* **83**:101733, 2023, doi:10.1016/j.intfin.2023.101733.

[31] R. Cont and P. Das, Rough volatility: Fact or artefact?, *Sankhyā: The Indian Journal of Statistics, Series B* **86** (2024), no. 1, 191–223, doi:10.1007/s13571-024-00322-2.

Email address: `vidal.llaurado@gmail.com`