

7520739-beyond-square-root-of-time-mixed-rough-volatility-liquidity-horizons-and-expecte

SUMMARY: This article examines how Expected Shortfall (ES) changes with the holding horizon when volatility combines a mean-reverting diffusive factor with a rough Volterra factor under the physical measure. The asset price remains a continuous semimartingale, so the leading short-horizon scale of returns and tail risk is proportional to $\sqrt{h}$; roughness enters through higher-order changes in the shape of the conditional loss distribution. For locally linear exposures, the paper derives explicit rough-leverage and integrated-variance corrections to Value-at-Risk (VaR) and ES. Under a weighted Fourier–Edgeworth condition, the rough-leverage correction is of relative order hH , whereas in the no-leverage variance-mixture regime, the leading symmetric correction is of order h^2H . The paper establishes perturbation results showing when these expansions survive pricing and linearization remainders, and develops fixed-horizon inference for the risk-functional parameters H and $\chi_R = \lambda R p_{SR}$, including a conditional delta method for normalized ES. A controlled Monte Carlo experiment verifies the sign, rate, and magnitude of the analytical correction and the finite-sample accuracy of the proposed standard errors. The framework is intended as a model-risk, economic-capital, and stress-testing tool, not as a replacement for regulatory liquidity-horizon rules.

METHOD: The study uses a mixed, mean-reverting Gaussian–Volterra specification in which a rough component describes short-scale fluctuations and a diffusive component allows a distinct medium- and long-horizon scale. The main theoretical point is that the traded price remains a continuous semimartingale, so the leading short-horizon return scale is still $\sqrt{h}$. The roughness enters through higher-order changes in the shape of the conditional loss distribution. The paper derives the leading local skewness and kurtosis generated by the Volterra volatility factors and obtains a uniform density and distribution expansion under an explicit weighted Fourier–Edgeworth regularity condition, which is then transferred to VaR and ES. The paper also develops an estimable representation of the leading correction, including transparent pilot estimators of the local Hurst exponent and the product $\chi_R = \lambda R p_{SR}$.

KEY FINDINGS:

- The rough-leverage correction to ES is of relative order hH .
- In the no-leverage variance-mixture regime, the leading symmetric correction to ES is of order h^2H .
- The paper establishes perturbation results showing when these expansions survive pricing and linearization remainders.
- A controlled Monte Carlo experiment verifies the sign, rate, and magnitude of the analytical correction and the finite-sample accuracy of the proposed standard errors.
- The framework is intended as a model-risk, economic-capital, and stress-testing tool, not as a replacement for regulatory liquidity-horizon rules.

IMPLICATION FOR TRADING: The findings of this paper have significant implications for risk management and trading practices. For traders and risk managers, understanding the impact of volatility and leverage on Expected Shortfall can help in making more informed decisions. The explicit corrections derived in the paper provide a more accurate understanding of risk measures under different conditions, which can be used to validate models and improve risk management practices. The framework

developed in the paper can serve as a tool for assessing and managing model risk, economic capital, and stress testing, which are crucial for financial institutions. However, it is important to note that this framework is not intended as a replacement for regulatory liquidity-horizon rules, and institutions should still adhere to these rules.

Beyond Square-Root-of-Time: Mixed Rough Volatility, Liquidity Horizons, and Expected Shortfall

Jules SADEFO KAMDEM

Montpellier Research in Economics (MRE, UR 209), Faculty of Economics,

University of Montpellier, France

jules.sadefo-kamdem@umontpellier.fr

ORCID: 0000-0002-7308-428X

September 2026

Abstract

We study how Expected Shortfall changes with the holding horizon when volatility combines a mean-reverting diffusive factor with a rough Volterra factor under the physical measure. The asset price remains a continuous semimartingale, so the leading short-horizon scale of returns and tail risk is proportional to $\sqrt{h}$; roughness enters through higher-order changes in the shape of the conditional loss distribution. For locally linear exposures, we derive explicit rough-leverage and integrated-variance corrections to Value-at-Risk and Expected Shortfall. Under a weighted Fourier–Edgeworth condition, the rough-leverage correction is of relative order h^H , whereas in the no-leverage variance-mixture regime the leading symmetric correction is of order h^{2H} . We establish perturbation results showing when these expansions survive pricing and linearization remainders, and develop fixed-horizon inference for the risk-functional parameters H and $\chi_R = \lambda_R \rho_{SR}$, including a conditional delta method for normalized ES. A controlled Monte Carlo experiment verifies the sign, rate and magnitude of the analytical correction and the finite-sample accuracy of the proposed standard errors. We then formulate a direct-horizon risk engine and a practitioner validation protocol for linear and nonlinear books. The framework is intended as a model-risk, economic-capital and stress-testing tool, not as a replacement for regulatory liquidity-horizon rules.

Keywords: Expected Shortfall; Value-at-Risk; rough volatility; Volterra processes; liquidity horizon; risk scaling; model risk; Monte Carlo.

JEL classification: C22; C58; G11; G17; G32.

1 Introduction

Market-risk measurement is intrinsically a multi-horizon problem. Trading positions are held, hedged or unwound over different periods, and the same institution may need a one-day risk signal for desk management, a ten-day measure for internal capital, and substantially longer

horizons for less liquid risk factors. A common operational shortcut is to transport a short-horizon volatility or risk estimate by a square-root-of-time factor. That rule is exact for sums of independent Gaussian increments with constant variance, but it need not be exact for conditional tail risk when volatility is stochastic, dependent across time and correlated with returns. The resulting model risk matters because Expected Shortfall (ES) is sensitive not only to conditional variance but also to the shape of the loss distribution in the tail.

The modern rough-volatility literature provides a natural setting in which to revisit this issue. Empirically, short-scale log-volatility fluctuations have often been found to be much less regular than a diffusion, motivating models with a Hurst parameter $H < 1/2$ (Gatheral et al., 2018; Bayer, P. Friz, et al., 2016). At the same time, evidence on roughness is not unambiguous: measurement error in realized volatility and scale-dependent persistence can materially affect the estimated short-run exponent (Bolko et al., 2023; Cont and Das, 2024; Bennedsen et al., 2022). This paper therefore does not assume that a single rough factor is the universally correct description of volatility. It uses a mixed, mean-reverting Gaussian–Volterra specification in which a rough component describes short-scale fluctuations and a diffusive component allows a distinct medium- and long-horizon scale.

The main theoretical point is deliberately narrower than the statement that “rough volatility breaks square-root-of-time.” The traded price remains a continuous semimartingale. Consequently, under standard local regularity conditions, the leading short-horizon return scale is still $\sqrt{h}$. Roughness appears in the higher-order shape of the conditional loss distribution. For a locally linear position, let $a_{t,h}$ denote the exact conditional standard-deviation scale of the martingale loss and let $c_\alpha = \phi(z_\alpha)/(1 - \alpha)$ be the Gaussian ES constant. Under a rough-leverage channel, the central expansion obtained below is

$$\frac{\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h}}{a_{t,h}c_\alpha} = 1 - \text{sgn}(\mathcal{E}_t) \frac{\chi_R z_\alpha}{2(H + \frac{1}{2})(H + \frac{3}{2})} h^H + o_p(h^H), \quad \chi_R = \lambda_R \rho_{SR}. \quad (1)$$

Thus the local diffusive scale is preserved, while rough volatility and return–volatility leverage generate an explicit finite-horizon correction. When return–volatility leverage is absent, uncertainty in integrated variance instead produces a symmetric excess-kurtosis correction of relative order h^{2H} . These two channels are economically distinct and need not have the same sign or relevance at a given holding period.

Obtaining an ES expansion from local cumulants requires more than a moment calculation. Because ES integrates the loss tail beyond the quantile, a pointwise central limit theorem or a Cornish–Fisher heuristic is insufficient. We therefore separate the argument into three steps. First, we derive the leading local skewness and kurtosis generated by the Volterra volatility factors. Second, under an explicit weighted Fourier–Edgeworth regularity condition, we obtain a uniform density and distribution expansion and transfer it to VaR and ES. Third, we prove stability of the risk-measure expansion when the martingale loss is perturbed by a dependent pricing or linearization remainder. The VaR result requires local anti-concentration and probability-tail control; the ES result follows from conditional L^1 stability and therefore needs only integrated tail control.

The paper also develops an estimable representation of the leading correction. The relevant

rough-leverage parameter enters the first-order ES term through $\chi_R = \lambda_R \rho_{SR}$. We construct transparent pilot estimators of the local Hurst exponent and this product, derive a fixed-horizon joint central limit theorem under a high-level influence-function condition, and propagate parameter uncertainty to the first-order ES approximation by a conditional delta method. A shrinking-horizon argument is presented only as an abstract structural interpretation: it shows how local effective parameters converge to primitive power-law parameters under additional infill conditions, but it is not used as a claim about a particular high-frequency sampling scheme.

The numerical analysis has two purposes. It does not constitute evidence that real-world volatility is rough. Instead, it verifies that the analytical coefficients have the correct rate and magnitude in a nonlinear rough-volatility simulation and that the proposed large-sample standard errors reproduce finite-sample dispersion in a controlled pilot-moment experiment. With $H = 0.12$, negative return–volatility correlation and 120,000 simulated paths, the first-order ES approximation remains within one percent of the nonlinear Monte Carlo ES over the reported one- to forty-day horizons. A second experiment with 5,000 replications shows close agreement between Monte Carlo standard deviations and the delta- method standard errors for H , χ_R and the ten-day ES shape multiplier.

The resulting quantities are intended for risk-management and actuarial use rather than as a proposed replacement for regulatory capital rules. Under the Basel internal-models approach, ES is computed at a 97.5% one-tailed confidence level and liquidity horizons are incorporated through a prescribed ten-day base-horizon aggregation rather than by scaling a one-day portfolio ES mechanically (Basel Committee on Banking Supervision, 2020). Our direct-horizon measures are therefore best interpreted as model-risk, economic-capital and stress-testing diagnostics. They can also be used in asset–liability management and hedging applications in which the financial asset portfolio is exposed to volatility-sensitive or nonlinear positions.

The paper contributes in three connected ways. Theoretically, it characterizes the first rough correction to VaR and ES while retaining the correct $\sqrt{h}$ leading scale, decomposes horizon transport into forward-variance and non-Gaussian shape effects, and extends the mechanism to nonlinear portfolios. Econometrically, it provides a risk- functional parameterization and asymptotic inference for the leading ES correction under the physical measure. Computationally, it validates the local formulas by Monte Carlo and specifies a reproducible direct-horizon engine for comparison with Gaussian, Student, GARCH, Markov stochastic-volatility and single-factor rough benchmarks. The empirical horse race is formulated as a validation protocol rather than reported as a result without data.

The remainder of the paper is organized as follows. Section 2 reviews the VaR/ES, temporal aggregation, fractional-volatility and rough-volatility literatures and positions the contribution. Section 3 introduces the mixed Gaussian–Volterra model. Section 4 derives the tail-risk asymptotics and perturbation results. Section 5 develops estimation and inference under the physical measure. Section 6 describes the direct-horizon risk engine and validation benchmarks. Section 7 reports the Monte Carlo validation. Section 8 specifies an implementable market and actuarial validation protocol, and Section 9 concludes.

2 Literature Review and Positioning

The paper lies at the intersection of four literatures that have largely developed in parallel: the theory and econometrics of Value-at-Risk (VaR) and Expected Shortfall (ES), the temporal aggregation of market risk, stochastic and fractional volatility modelling, and the recent rough-volatility literature. The central issue is not whether the square-root-of-time rule is “wrong” in every setting. For a continuous semimartingale price with locally finite spot variance, the leading short-horizon scale of returns remains proportional to the square root of the horizon. The economically relevant question is instead how rapidly the *finite-horizon* tail of portfolio losses departs from that leading-order approximation when volatility is random, strongly time dependent, rough at short scales, persistent at longer scales, and correlated with returns. This distinction is important for the interpretation of the title of the paper: “beyond square-root-of-time” refers to model-implied horizon corrections at holding periods relevant for risk management, rather than to a blanket replacement of the local diffusive rate by h^H .

2.1 From static VaR to conditional Expected Shortfall

VaR entered modern risk management as a quantile-based summary of the loss distribution, but its limitations as a general measure of portfolio risk motivated the axiomatic theory of coherent risk measures. In particular, Artzner et al. (1999) formalized monotonicity, translation invariance, positive homogeneity and subadditivity, while Acerbi and Tasche (2002) clarified the coherence of ES and the care required when loss distributions are not continuous. The optimization representation of Conditional Value-at-Risk in Rockafellar and Uryasev (2000) also made tail-mean risk operational for portfolio construction and scenario-based optimization.

A complementary strand of research has focused on obtaining VaR and ES analytically under non-Gaussian portfolio distributions. This is directly relevant to the present paper because rough stochastic volatility ultimately produces non-Gaussian unconditional loss distributions even when returns are conditionally Gaussian. Earlier analytical work by Sadefo Kamdem (2005) derives VaR and ES for linear portfolios with elliptically distributed risk factors, including the multivariate Student case, while Sadefo Kamdem (2012) extends the analysis to mixtures of generalized Laplace distributions with time-varying tail shape. For nonlinear portfolios, Sadefo Kamdem and Genz (2008) develops numerical integration methods for quadratic portfolios with options under multivariate elliptical risk factors. These contributions illustrate a general principle that remains central here: tail-risk measures depend not only on conditional variance but also on the entire geometry of the portfolio loss distribution, including heavy tails, dependence and nonlinear exposures.

The econometric literature subsequently shifted from unconditional or one-period distributional formulas to conditional tail forecasting. McNeil and Frey (2000) combine GARCH-type volatility filtering with extreme-value methods for conditional VaR and ES. Engle and Manganelli (2004) model conditional quantiles directly through CAViaR, while Cai and Wang (2008) propose nonparametric estimators of conditional VaR and ES. These approaches build on the broader conditional-heteroskedasticity tradition initiated by Engle (1982) and Bollerslev (1986). Forecast evaluation is equally important: Christoffersen (1998) provides conditional coverage tests for

interval forecasts, and Du and Escanciano (2017) develops ES backtests that account for the severity of tail losses.

A major statistical advance is the joint elicibility of VaR and ES established by Fissler and Ziegel (2016). This result makes it possible to compare competing VaR–ES forecasts using strictly consistent joint loss functions rather than relying only on exception counts. Patton et al. (2019) exploit this structure to construct dynamic semiparametric models of VaR and ES and document competitive out-of-sample performance. For the present paper, these results imply that a mixed-rough model must not be judged merely by likelihood fit or by the magnitude of its ES. It must be evaluated as a *tail-forecasting model*, using proper joint VaR–ES scores and formal backtests.

The contribution pursued here therefore differs from the classical distributional VaR/ES problem. The object of interest is a family of conditional loss distributions $\{L_{t,t+h} : h > 0\}$ generated by one structural volatility mechanism. Instead of estimating a separate risk number at each frequency, we ask whether the same stochastic-volatility law can explain the *term structure* of tail risk across economically relevant holding periods.

2.2 Temporal aggregation, square-root-of-time scaling and liquidity horizons

The square-root-of-time rule follows exactly for centered sums of independent Gaussian increments with constant variance. Its practical attraction is obvious: a one-day risk estimate can be transported to a longer horizon by multiplying the volatility component by $\sqrt{h}$. Yet temporal aggregation is not innocuous once volatility is time varying. Drost and Nijman (1993) show that aggregation changes the parameters of GARCH processes and that the common assumption that rescaled innovations remain independent is generally questionable. For tail quantiles, Daniélsson and Zigrand (2006) demonstrate in a jump-diffusion setting that square-root scaling can systematically underestimate risk when jumps are losses, with the discrepancy increasing with the horizon, jump intensity and confidence level.

The recent literature has moved from rejecting a universal scaling rule to constructing data-driven risk-scaling procedures. Pitera et al. (2026) develop an adjustment framework for risk estimators that explicitly accommodates transport across horizons and confidence levels and links the scaling correction to backtesting performance. This is particularly relevant because it shows that an empirical scaling coefficient or exponent, by itself, is no longer a sufficient novelty. A structural contribution must explain *why* the correction varies with horizon and market state.

Liquidity-horizon risk constitutes a related but distinct strand. When the liquidation period is itself uncertain, the distribution of portfolio losses becomes a mixture over holding periods. Colldeforms-Papiol and Ortiz-Gracia (2018) compute VaR and ES under stochastic liquidity horizons using delta–gamma portfolio approximations and jump models, illustrating that horizon uncertainty and payoff nonlinearity interact in a nontrivial way.

The regulatory treatment is also more subtle than the shorthand expression $ES_h = \sqrt{h} ES_1$. Under the current Basel internal-models approach, Basel Committee on Banking Supervision (2020), ES is computed daily at a one-tailed 97.5% confidence level. The base horizon is ten days, with prescribed liquidity horizons of 10, 20, 40, 60 and 120 days. Importantly, the ten-day quantities entering the liquidity-horizon aggregation must be based on risk-factor changes

over ten days rather than obtained by scaling a shorter- horizon ES. Square-root factors enter subsequently through the regulatory aggregation of nested risk-factor sets assigned to different liquidity horizons. Therefore, a useful model-risk study should not compare a rough-volatility ES to a fictitious “one-day FRTB ES times $\sqrt{10}$ ”. The economically meaningful comparison is between (i) direct model-implied loss distributions at each horizon and (ii) the risk obtained from the regulatory or conventional horizon-transport architecture.

This distinction motivates one of the paper’s key objects. For a reference horizon h_0 , define the horizon-scaling distortion

$$\mathcal{D}_{\alpha,t}(h \mid h_0) = \frac{ES_{\alpha,t}(h)}{ES_{\alpha,t}(h_0)\sqrt{h/h_0}} - 1. \quad (2)$$

The statistic is descriptive rather than regulatory: $\mathcal{D}_{\alpha,t} = 0$ means that the model-implied ES is consistent with square-root transport from h_0 ; positive and negative values quantify, respectively, under- and over-scaling by that benchmark. The novel issue is to link the sign and magnitude of (2) to primitive features of volatility dynamics rather than to estimate an unrestricted empirical exponent.

2.3 From fractional long memory to rough stochastic volatility

Fractional Brownian motion (fBm), introduced in its modern moving-average form by Mandelbrot and Van Ness (1968), provides a parsimonious way to generate non-Markovian dependence through the Hurst parameter H . In financial volatility modelling, the first major continuous-time applications emphasized persistence and long memory. The fractional stochastic-volatility model of Comte and Renault (1998), for example, was constructed in the regime associated with persistent volatility. Mixed fractional models also combine Brownian and fractional components. However, when such a mixture is placed directly in the asset-price dynamics, semimartingale issues become central: Cheridito (2001) characterizes when a Brownian motion plus an independent fBm is a semimartingale. This is one reason the present paper puts the non-Markovian component in *volatility* while retaining a standard semimartingale price.

There is also a direct connection with earlier mixed-fractional work on option valuation. Djetcha and Sadefo Kamdem (2021) derive local and implied volatility formulas in a mixed-modified-fractional Dupire framework and study their dependence on the Hurst parameter. That line of work concerns a long-memory mixed-fractional pricing environment. The present paper makes a different modelling choice: it uses a rough component ($H < 1/2$) in the volatility under the physical measure and asks how that component affects the horizon profile of portfolio losses. The distinction between long-memory fractional pricing models and modern rough volatility is therefore substantive, not terminological.

Against the classical Markovian stochastic-volatility benchmark of Heston (1993), the modern rough-volatility literature begins from the empirical observation that log-realized volatility is much less regular than a diffusion. Gatheral et al. (2018) estimate Hurst exponents of order 0.1 in major equity data and propose the rough fractional stochastic-volatility paradigm. In pricing, Bayer, P. Friz, et al. (2016) introduce the rough Bergomi model, while El Euch and Rosenbaum (2019) obtain a fractional-Riccati representation for the characteristic function of rough Heston.

The affine-Volterra framework of Abi Jaber et al. (2019) places such models in a broader class of stochastic convolution equations and supplies transform techniques that are valuable for both pricing and risk computation.

Short-time asymptotics provide a second theoretical pillar. Forde and Zhang (2017) derive large-deviation asymptotics for rough stochastic-volatility models, and Bayer, P. K. Friz, et al. (2019) develop moderate-deviation expansions near the money. These results show how fractional regularity enters distributional corrections and implied-volatility asymptotics without implying that the asset return itself scales as h^H . This point is essential for tail risk. If

$$dX_t = \left(\mu_t - \frac{1}{2} v_t \right) dt + \sqrt{v_t} dW_t, \quad (3)$$

and v_t is locally well behaved in the sense required for $\int_t^{t+h} v_s ds = v_t h + o_p(h)$, then

$$\frac{X_{t+h} - X_t}{\sqrt{h}} \implies \sqrt{v_t} Z, \quad h \downarrow 0, \quad (4)$$

conditionally on current information, with Z standard normal in the simplest case. Thus rough volatility need not overturn the leading $\sqrt{h}$ return scale. Its effect can appear in higher-order distributional terms, through random integrated variance, leverage, volatility-of-volatility and nonlinear portfolio revaluation. This observation changes the theoretical target from “estimating a new universal exponent” to deriving the *rough correction to finite-horizon ES*.

2.4 Is volatility rough, persistent, or multi-scale?

The empirical status of roughness is now sufficiently nuanced that it should be treated as a testable feature rather than imposed as a stylized fact. Measurement error is especially important because spot volatility is latent. Bolko et al. (2023) develop a GMM estimator for fractional log-normal stochastic volatility and show analytically that noisy realized measures can bias estimated roughness; after correcting for this effect, they still obtain very low Hurst estimates on major equity indices. By contrast, Cont and Das (2024) show through a model-free analysis and simulation experiments that realized volatility can appear rough even when the underlying instantaneous volatility is diffusive, strengthening the case for explicit measurement-error controls. Option-market arguments provide an alternative source of information: Fukasawa (2021) shows that, under a power-law short-maturity at-the-money skew, sufficiently regular volatility dynamics are incompatible with a viable continuous-price market, thereby giving a theoretical motivation for rough volatility from option smiles rather than realized-volatility scaling alone.

A second issue is that one fractional exponent need not describe all horizons. Bennedsen et al. (2022) construct stochastic-volatility models that decouple short-run roughness from slowly decaying long-run dependence and show substantial volatility- forecasting gains from allowing both features. Garcin and Grasselli (2022) similarly document scale-dependent behavior: very low apparent Hurst exponents at short scales coexist with a different geometry at longer scales. These findings are directly aligned with the present paper’s multi-horizon objective. If the short end of the volatility process and its medium-run persistence are governed by different

components, the term structure of ES should inherit a crossover rather than obey one global power law.

This motivates a mixed, mean-reverting Volterra specification rather than a raw sum of a Brownian motion and a nonstationary fBm in log volatility. Schematically, the empirical model considered later in the paper has

$$\log v_t = m + X_t^D + X_t^R, \quad (5)$$

where X^D is a diffusive mean-reverting component and X^R is a rough Volterra component whose kernel behaves like $u^{H-1/2}$ near the origin and is tempered or mean reverting at longer lags. The simulation results of Bennedsen et al. (2017) make this class computationally feasible, while the empirical results above provide a reason to separate short- and long-scale behavior.

The adjective “mixed” must also be used carefully. It is not new in the rough-volatility literature. Bourgey et al. (2023), building on earlier constructions in the VIX literature, explicitly use the term *mixed rough Bergomi* for an enhanced forward- variance model that improves VIX-smile flexibility. Accordingly, the novelty claimed here is not the invention of a mixed rough-volatility model as such. It is the use of a parsimonious mixed-rough mechanism under the physical measure to explain and forecast the *horizon dependence of portfolio tail risk*.

2.5 Rough volatility, tail risk, and the remaining gap

Despite the rapid development of rough-volatility pricing, comparatively few papers study portfolio tail risk directly. Morelli and Santucci de Magistris (2019) analyze the tail behavior of fractional Ornstein–Uhlenbeck volatility and the probability that volatility exceeds high thresholds. Their object is primarily *volatility tail risk*, not the multi-horizon loss distribution of a general portfolio. Dupret and Hainaut (2021) provide a much closer precedent: they study CPPI strategies under affine Volterra processes, including rough Heston, and compute VaR and tail-mean risk measures using the characteristic function of the log cushion. This paper is important for delimiting our novelty: Expected Shortfall under rough Heston is not new, and roughness does not mechanically imply a larger ES for every parameterization.

The remaining issue is therefore not whether an ES number can be computed under a rough model. It is whether a structural rough/multi-scale volatility law can explain the *cross-horizon transformation* of the loss distribution and whether that transformation has incremental forecasting value relative to standard alternatives. This question becomes especially important for nonlinear portfolios. For an option book, a local approximation of the profit and loss contains terms of the form

$$\Delta V \approx \Delta \Delta S + \frac{1}{2} \Gamma (\Delta S)^2 + \text{Vega} \Delta \sigma + \text{Vanna} \Delta S \Delta \sigma + \frac{1}{2} \text{Volga} (\Delta \sigma)^2. \quad (6)$$

Even when the leading Delta component inherits the local diffusive $\sqrt{h}$ rate, the Vega, Vanna and Volga terms expose the portfolio directly to rough volatility increments and to the leverage channel linking ΔS and $\Delta \sigma$. Hence a small rough-volatility effect on a cash-equity ES can coexist with a much larger effect on the ES of an options book. This also creates a natural bridge to the earlier nonlinear VaR literature represented by Sadefo Kamdem and Genz (2008).

A further distinction concerns probability measures. Most rough-volatility pricing models are calibrated under the risk-neutral measure $\mathbb{Q}$, whereas market-risk VaR and ES are predictive objects under the physical measure $\mathbb{P}$. Transporting a risk-neutral rough calibration directly into a historical P&L generator would therefore confound physical dynamics with the volatility risk premium. Bonesini et al. (2025) make this distinction explicit by starting from a rough-volatility model under the historical measure and studying the effect of a stochastic volatility risk premium on pricing. Our empirical design follows the same separation of roles: the mixed-rough dynamics used for tail forecasting are estimated under $\mathbb{P}$, while $\mathbb{Q}$ enters only where conditional derivative revaluation requires it.

Taken together, the literature has developed, separately, coherent ES, dynamic VaR–ES forecasting and backtesting, risk scaling, regulatory liquidity horizons, rough and multi-scale volatility, and selected rough-volatility tail-risk applications. We are not aware of a systematic study that combines all of the following elements in one framework: (i) a mixed rough/multi-scale volatility mechanism estimated under the physical measure, (ii) direct ES forecasts over a grid of holding periods aligned with practical liquidity horizons, (iii) an explicit decomposition of the deviation from square-root transport, (iv) nonlinear option-book revaluation, and (v) rolling out-of-sample joint VaR–ES forecast evaluation against both non-rough and rough benchmarks. This is the gap targeted here.

2.6 Positioning and contributions

The paper is positioned as a model-risk and economic-risk study, not as a proposal to replace the Basel market-risk formula. Square-root transport is treated as a leading-order benchmark, the Basel liquidity-horizon construction as a regulatory benchmark, and direct- horizon simulation under the mixed-rough model as a structural benchmark. The purpose of the estimation and validation framework is to determine when the additional rough and multi-scale structure has measurable value for tail-risk forecasting rather than to assume that it does.

The paper makes three theoretical contributions:

- T1. Finite-horizon rough correction to VaR and ES.** We preserve the $\sqrt{h}$ leading scale implied by the semimartingale price process and identify explicit higher-order corrections generated by rough leverage and random integrated variance. Weighted Edgeworth control is separated from cumulant calculations, and perturbation results state when the martingale-loss expansion transfers to the full P&L.
- T2. Structural horizon distortion and crossover.** We decompose horizon transport into a forward-variance scale effect and a non-Gaussian shape effect, and relate the latter to H , the rough amplitude and return–volatility leverage. This makes the Rough Horizon Multiplier a model-risk diagnostic rather than an unrestricted empirical scaling exponent.
- T3. Nonlinear portfolio transmission.** We show how volatility-sensitive nonlinear exposures create additional channels through Delta, Gamma, Vega, Vanna and Volga. The framework therefore distinguishes the local theory for linear losses from the full-revaluation problem relevant to derivatives books.

The paper makes three methodological and computational contributions:

- M1. **Physical-measure risk inference.** We formulate measurement-error-aware pilot moments for the short-scale volatility geometry, derive fixed-horizon asymptotic inference for the effective parameters (H, χ_R) with $\chi_R = \lambda_R \rho_{SR}$, and propagate their uncertainty to the first-order ES correction.
- M2. **Direct-horizon Monte Carlo validation.** We verify the sign, rate and magnitude of the local rough-leverage correction in a nonlinear simulation and assess finite-sample recovery of H , χ_R and the ES multiplier. The risk engine is constructed to simulate the target holding horizon directly rather than to infer it from a shorter-horizon risk number.
- M3. **A reproducible professional validation protocol.** We specify how the model can be evaluated against Gaussian, Student, GARCH- t , Markov stochastic-volatility and single-factor rough alternatives using VaR calibration, ES backtests, joint VaR–ES scoring and economic materiality. The protocol covers linear and nonlinear books and separates economic/model-risk diagnostics from the regulatory Basel liquidity-horizon calculation.

The resulting research question is deliberately operational: *does a rough and multi-scale volatility mechanism explain enough of the horizon dependence of portfolio tail losses to improve risk measurement after accounting for estimation, simulation and model uncertainty?*

3 Mixed Rough Volatility and Horizon-Dependent Tail Risk

This section introduces the structural model used to study how portfolio tail risk evolves with the holding horizon. The construction is deliberately designed around three principles. First, the traded asset price remains a continuous semimartingale under the physical measure, so that the model is compatible with the standard local no-arbitrage architecture of continuous-time finance. Second, the volatility process is allowed to exhibit genuinely rough behavior at short scales while retaining a distinct mean-reverting component at medium and longer scales. Third, the model is formulated under the physical probability measure because VaR and Expected Shortfall are predictive risk measures. Risk-neutral quantities enter only when a derivative portfolio has to be revalued inside a physical-measure scenario.

The resulting specification is a mixed Gaussian–Volterra log-volatility model. It is not intended to claim novelty for the use of multiple rough or Volterra factors per se; such constructions already appear in the rough-volatility literature. The innovation here is the role assigned to the mixture: it generates a structural, horizon-dependent distortion of the portfolio loss distribution that can be compared directly with conventional and regulatory risk transport across holding periods.

3.1 Filtered probability space and physical-measure price dynamics

Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a filtered probability space satisfying the usual conditions. The filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ supports a three-dimensional Brownian motion

$$\mathbf{B}_t = (W_t^S, B_t^D, B_t^R)',$$

with instantaneous correlation matrix

$$\mathbf{R} = \begin{pmatrix} 1 & \rho_{SD} & \rho_{SR} \\ \rho_{SD} & 1 & \rho_{DR} \\ \rho_{SR} & \rho_{DR} & 1 \end{pmatrix}, \quad \mathbf{R} \succeq 0. \quad (7)$$

The superscripts D and R denote, respectively, a diffusive mean-reverting volatility factor and a rough Volterra factor. The correlations ρ_{SD} and ρ_{SR} generate return–volatility leverage, while ρ_{DR} allows the two volatility components to share innovations. For the model definition, positive semidefiniteness is sufficient. The density and Edgeworth arguments in Section 4 work on a compact subregion on which the correlation matrix is uniformly positive definite; this excludes singular return–volatility configurations without changing the economic specification.

Under the physical measure, the equity index or tradable asset follows

$$\frac{dS_t}{S_t} = \mu_t dt + \sqrt{v_t} dW_t^S, \quad S_0 > 0, \quad (8)$$

where μ_t is locally bounded and predictable and $v_t > 0$ is the instantaneous variance. Equivalently, for $X_t = \log S_t$,

$$dX_t = \left(\mu_t - \frac{1}{2}v_t \right) dt + \sqrt{v_t} dW_t^S. \quad (9)$$

The semimartingale structure is entirely carried by the price equation. The volatility factor itself need not be a semimartingale. This separation is important when $H < 1/2$: rather than inserting a fractional Brownian motion directly into the tradable price, as in some mixed-fractional pricing models, roughness enters through the stochastic environment in which the semimartingale price evolves. This is the standard conceptual separation behind modern rough-volatility models such as rough Bergomi and rough Heston (Bayer, P. Friz, et al., 2016; El Euch and Rosenbaum, 2019).

3.2 A stationary mixed rough log-volatility specification

We model log variance as the sum of two stationary Gaussian Volterra factors,

$$\log v_t = \log \bar{v} + Y_t - \frac{1}{2}\omega_Y^2, \quad Y_t = \sigma_D X_t^D + \sigma_R X_t^R, \quad (10)$$

where $\bar{v} > 0$ is the unconditional mean of v_t and $\omega_Y^2 = \text{Var}(Y_t)$. The variance correction in (10) ensures $\mathbb{E}[v_t] = \bar{v}$ under stationarity. The two normalized factors are

$$X_t^D = \int_{-\infty}^t K_D(t-s) dB_s^D, \quad K_D(u) = \sqrt{2\kappa_D} e^{-\kappa_D u} \mathbf{1}_{\{u>0\}}, \quad (11)$$

$$X_t^R = \int_{-\infty}^t K_R(t-s) dB_s^R, \quad K_R(u) = c_{H,\kappa_R} u^{H-\frac{1}{2}} e^{-\kappa_R u} \mathbf{1}_{\{u>0\}}, \quad (12)$$

with

$$c_{H,\kappa_R} = \frac{(2\kappa_R)^H}{\sqrt{\Gamma(2H)}}, \quad 0 < H < \frac{1}{2}, \quad \kappa_D, \kappa_R > 0. \quad (13)$$

The normalization is chosen so that

$$\int_0^\infty K_D(u)^2 du = \int_0^\infty K_R(u)^2 du = 1, \quad (14)$$

which makes σ_D and σ_R directly interpretable as unconditional volatility- factor scales.

The diffusive component X^D is an Ornstein–Uhlenbeck factor with covariance

$$\text{Cov}(X_t^D, X_{t+h}^D) = e^{-\kappa_D|h|}. \quad (15)$$

The rough component is a tempered Brownian semistationary process. Near the origin its kernel behaves as $u^{H-1/2}$, generating rough increments of order h^H , while the exponential tempering $e^{-\kappa_R u}$ prevents the power law from governing all horizons. This separation is consistent with the empirical motivation for decoupling short- and long-scale volatility behavior (Bennedsen et al., 2022; Garcin and Grasselli, 2022).

Because the Brownian drivers B^D and B^R may be correlated, the unconditional variance of Y_t is

$$\omega_Y^2 = \sigma_D^2 + \sigma_R^2 + 2\rho_{DR}\sigma_D\sigma_R J_{DR}, \quad (16)$$

where

$$J_{DR} = \int_0^\infty K_D(u)K_R(u)du = \frac{\sqrt{2\kappa_D}(2\kappa_R)^H \Gamma\left(H + \frac{1}{2}\right)}{\sqrt{\Gamma(2H)}(\kappa_D + \kappa_R)^{H+\frac{1}{2}}}. \quad (17)$$

Equation (17) provides a convenient scalar measure of the overlap between the diffusive and rough kernels. In particular, J_{DR} is close to zero when the two factors operate at sharply different mean-reversion scales, and is larger when their frequency ranges overlap.

Remark 3.1 (Alternative baseline variance). A deterministic seasonal or slowly varying baseline $\bar{v}_t$ can replace $\bar{v}$ in (10). For derivative-pricing applications it is also useful to rewrite the model conditionally in terms of the forward variance curve $\xi_t(u) = \mathbb{E}_t[v_{t+u}]$, as in rough Bergomi. The stationary formulation is retained as the baseline under $\mathbb{P}$ because it makes the multi-horizon identification problem transparent and separates long-run mean reversion from short-run roughness.

3.3 Short-scale roughness and the volatility crossover

The central primitive of the model is the coexistence of two local regularities. The Ornstein–Uhlenbeck component has increments of standard deviation proportional to $h^{1/2}$, while the Volterra component has increments of standard deviation proportional to h^H . Since $H < 1/2$, the rough component dominates sufficiently short changes in log variance even when its unconditional variance contribution is modest.

Define the log-volatility variogram

$$\mathcal{V}_Y(h) = \mathbb{E}[(Y_{t+h} - Y_t)^2]. \quad (18)$$

For the diffusive factor,

$$\mathbb{E}[(X_{t+h}^D - X_t^D)^2] = 2(1 - e^{-\kappa_D h}) = 2\kappa_D h + o(h). \quad (19)$$

For the rough factor, the local singularity of K_R implies

$$\mathbb{E}[(X_{t+h}^R - X_t^R)^2] = C_R(H, \kappa_R) h^{2H} + o(h^{2H}), \quad (20)$$

where

$$C_R(H, \kappa_R) = c_{H, \kappa_R}^2 \left\{ \int_0^\infty \left[(x+1)^{H-\frac{1}{2}} - x^{H-\frac{1}{2}} \right]^2 dx + \frac{1}{2H} \right\} > 0. \quad (21)$$

The exponential tempering is asymptotically negligible at the scale $h \downarrow 0$ but enters the coefficient through the normalization c_{H, κ_R} .

The cross-factor increment covariance is of intermediate order,

$$\text{Cov}(X_{t+h}^D - X_t^D, X_{t+h}^R - X_t^R) = C_{DR}(H, \kappa_D, \kappa_R) h^{H+\frac{1}{2}} + o(h^{H+\frac{1}{2}}), \quad (22)$$

for a finite constant C_{DR} . Combining (19)–(22) yields

$$\mathcal{V}_Y(h) = \sigma_R^2 C_R h^{2H} + 2\rho_{DR} \sigma_D \sigma_R C_{DR} h^{H+\frac{1}{2}} + 2\kappa_D \sigma_D^2 h + o(h). \quad (23)$$

Because

$$2H < H + \frac{1}{2} < 1,$$

the rough term is the leading contribution as $h \downarrow 0$ whenever $\sigma_R > 0$. The diffusive factor can nevertheless dominate over intermediate or long horizons because the rough component is tempered and mean reverting.

When $\rho_{DR} = 0$, a useful primitive diagnostic is the rough–diffusive crossover obtained by equating the two leading variogram contributions,

$$\sigma_R^2 C_R (h^*)^{2H} = 2\kappa_D \sigma_D^2 h^*. \quad (24)$$

Hence

$$h^* = \left(\frac{\sigma_R^2 C_R(H, \kappa_R)}{2\kappa_D \sigma_D^2} \right)^{\frac{1}{1-2H}} \quad (25)$$

provided $\sigma_D \sigma_R > 0$. We refer to h^* as the *volatility crossover horizon*. It is not itself a VaR or ES horizon. Rather, it identifies the time scale around which the source of volatility variation changes, which in turn can generate a crossover in tail-risk scaling.

Remark 3.2 (Cross-factor correlation). When $\rho_{DR} \neq 0$, the intermediate term in (23) can shift the empirical crossover. In the empirical section we therefore report both the primitive benchmark h^* in (25) and an effective crossover obtained numerically from the full variogram. This prevents the scalar h^* from being overinterpreted when the two volatility shocks are strongly correlated.

3.4 Integrated variance and the conditional law of returns

For a holding horizon $h > 0$, define integrated variance and integrated drift by

$$IV_{t,h} = \int_t^{t+h} v_s ds, \quad A_{t,h} = \int_t^{t+h} \left(\mu_s - \frac{1}{2} v_s \right) ds. \quad (26)$$

The log return is

$$R_{t,h} := X_{t+h} - X_t = A_{t,h} + \int_t^{t+h} \sqrt{v_s} dW_s^S. \quad (27)$$

If the return Brownian motion is independent of the volatility drivers, $\rho_{SD} = \rho_{SR} = 0$, then conditional on the variance path,

$$R_{t,h} \mid \{v_s : t \leq s \leq t+h\} \sim \mathcal{N}(A_{t,h}, IV_{t,h}). \quad (28)$$

Thus the unconditional or $\mathcal{F}_t$ -conditional return law is a normal variance mixture. Randomness in $IV_{t,h}$ generates excess kurtosis even when the instantaneous innovation is Gaussian.

With leverage, (28) no longer holds in this simple form because the same Gaussian shocks influence both the future variance path and the return. A useful representation is obtained by collecting the two volatility Brownian motions into $d\mathbf{B}_t^v = (dB_t^D, dB_t^R)'$ with correlation matrix

$$\mathbf{R}_v = \begin{pmatrix} 1 & \rho_{DR} \\ \rho_{DR} & 1 \end{pmatrix}, \quad \mathbf{r} = \begin{pmatrix} \rho_{SD} \\ \rho_{SR} \end{pmatrix}. \quad (29)$$

When $\mathbf{R}_v$ is nonsingular, positive semidefiniteness of (7) implies

$$q_S := 1 - \mathbf{r}' \mathbf{R}_v^{-1} \mathbf{r} \geq 0. \quad (30)$$

We may therefore write

$$dW_t^S = \mathbf{r}' \mathbf{R}_v^{-1} d\mathbf{B}_t^v + \sqrt{q_S} dW_t^\perp, \quad (31)$$

where $W^\perp$ is independent of the volatility drivers. Substituting into (27),

$$R_{t,h} = A_{t,h} + \int_t^{t+h} \sqrt{v_s} \mathbf{r}' \mathbf{R}_v^{-1} d\mathbf{B}_s^v + \sqrt{q_S} \int_t^{t+h} \sqrt{v_s} dW_s^\perp. \quad (32)$$

Conditionally on the volatility-driver paths, only the last term remains Gaussian. This representation is useful both conceptually and computationally: it isolates the leverage contribution to skewness from the residual Gaussian variation.

3.5 Direct-horizon loss distributions

Let Π_t denote the marked-to-market value of a portfolio. We define the loss over $[t, t+h]$ as

$$L_{t,h} = \Pi_t - \Pi_{t+h}. \quad (33)$$

All risk measures in the paper are computed from the conditional distribution of $L_{t,h}$ given $\mathcal{F}_t$. This *direct-horizon* definition is central: the h -day distribution is generated from the model over

h days rather than obtained by rescaling a one-day risk number.

For a linear cash position of n shares,

$$L_{t,h}^{\text{cash}} = -n(S_{t+h} - S_t). \quad (34)$$

For short horizons it is often convenient to use the local approximation

$$L_{t,h}^{\text{cash}} = -nS_t R_{t,h} + O_p(h), \quad (35)$$

which is sufficient for the local tail expansions derived in Section 4. The direct Monte Carlo calculations later in the paper use the exact price change rather than (35).

For a derivatives book, let

$$\Pi_t = \mathcal{P}(t, \mathcal{X}_t; \Theta_t^{\mathbb{Q}}), \quad (36)$$

where $\mathcal{X}_t$ denotes the physical state variables observed at time t and $\Theta_t^{\mathbb{Q}}$ denotes the pricing parameters or implied-volatility state required for risk-neutral valuation. A physical-measure scenario first evolves the market state to $t+h$ and then revalues the portfolio using the pricing map at that horizon. This produces

$$L_{t,h}^{\text{full}} = \mathcal{P}(t, \mathcal{X}_t; \Theta_t^{\mathbb{Q}}) - \mathcal{P}(t+h, \mathcal{X}_{t+h}; \Theta_{t+h}^{\mathbb{Q}}). \quad (37)$$

The corresponding Greek approximation is

$$\Delta\Pi \approx \Delta\Delta S + \frac{1}{2}\Gamma(\Delta S)^2 + \text{Vega}\Delta\sigma + \text{Vanna}\Delta S\Delta\sigma + \frac{1}{2}\text{Volga}(\Delta\sigma)^2 + \Theta h, \quad (38)$$

with $L \approx -\Delta\Pi$. Equation (38) is an operational approximation for attribution and fast risk calculation; the full-revaluation loss in (37) is the benchmark for nonlinear books.

Remark 3.3 (No contradiction with semimartingale pricing). A rough volatility factor can move by order h^H , but this should not be interpreted as a claim that a correctly priced traded option has a universal h^H P&L scaling. Discounted traded-claim prices remain semimartingales under an equivalent risk-neutral measure. The Greek decomposition is used here to identify finite-horizon channels—especially Vega, Vanna and Volga—through which rough volatility affects the loss distribution, not to replace the semimartingale small-time limit of the traded claim.

3.6 Conditional VaR, Expected Shortfall and the horizon term structure

For confidence level $\alpha \in (0, 1)$, define the conditional distribution function

$$F_{t,h}(\ell) = \mathbb{P}(L_{t,h} \leq \ell \mid \mathcal{F}_t). \quad (39)$$

The direct-horizon VaR is

$$\text{VaR}_{\alpha,t}(h) = \inf\{\ell : F_{t,h}(\ell) \geq \alpha\}. \quad (40)$$

For a continuous conditional loss distribution, Expected Shortfall is

$$\text{ES}_{\alpha,t}(h) = \mathbb{E}[L_{t,h} \mid L_{t,h} \geq \text{VaR}_{\alpha,t}(h), \mathcal{F}_t], \quad (41)$$

with the equivalent quantile representation

$$\text{ES}_{\alpha,t}(h) = \frac{1}{1-\alpha} \int_{\alpha}^1 \text{VaR}_{u,t}(h) du. \quad (42)$$

The collection

$$\mathcal{T}_{\alpha,t}^{ES} = \{\text{ES}_{\alpha,t}(h) : h \in \mathcal{H}\} \quad (43)$$

for a set of holding periods $\mathcal{H}$ is the *Expected-Shortfall term structure*. In the empirical work we use horizons that include 1, 5 and 10 trading days as well as the 10, 20, 40, 60 and 120 day liquidity-horizon grid used in market-risk practice.

Given a reference horizon h_0 , define the square-root transport benchmark

$$\text{ES}_{\alpha,t}^{\sqrt{\cdot}}(h | h_0) = \text{ES}_{\alpha,t}(h_0) \sqrt{\frac{h}{h_0}}. \quad (44)$$

The corresponding *Rough Horizon Multiplier* is

$$\text{RHM}_{\alpha,t}(h | h_0) = \frac{\text{ES}_{\alpha,t}(h)}{\text{ES}_{\alpha,t}(h_0) \sqrt{h/h_0}}, \quad (45)$$

with scaling distortion

$$\mathcal{D}_{\alpha,t}(h | h_0) = \text{RHM}_{\alpha,t}(h | h_0) - 1. \quad (46)$$

Unlike an unrestricted empirical power exponent, the purpose of (45) is to connect horizon distortion to the primitive volatility parameters $(H, \sigma_R, \sigma_D, \kappa_R, \kappa_D, \rho_{SR}, \rho_{SD})$ and to the portfolio's nonlinear exposures.

A local elasticity summarizes the slope of the ES term structure,

$$\beta_{\alpha,t}(h) = \frac{\partial \log \text{ES}_{\alpha,t}(h)}{\partial \log h}, \quad (47)$$

whenever the derivative exists. The benchmark value 1/2 is interpreted as a local reference, not as a maintained restriction over all horizons.

3.7 Why square-root-of-time remains the local benchmark

The roughness of v does not by itself imply that the return in (9) scales as h^H . The next proposition formalizes the local benchmark that motivates the higher-order analysis in the next section.

Assumption 3.4 (Local continuity). *At the risk-measurement date t , $v_t > 0$ almost surely, $v_s \rightarrow v_t$ in probability as $s \downarrow t$, and μ is locally bounded. The portfolio is locally linear in the underlying with dollar exposure $\mathcal{E}_t \neq 0$.*

Proposition 3.5 (Leading short-horizon VaR and ES). *Under Assumption 3.4, for the locally linear loss $L_{t,h} = -\mathcal{E}_t(X_{t+h} - X_t) + o_p(\sqrt{h})$ and for $\alpha > 1/2$,*

$$\text{VaR}_{\alpha,t}(h) = |\mathcal{E}_t| \sqrt{v_t} z_{\alpha} \sqrt{h} + o_p(\sqrt{h}), \quad (48)$$

$$\text{ES}_{\alpha,t}(h) = |\mathcal{E}_t| \sqrt{v_t} \frac{\phi(z_\alpha)}{1-\alpha} \sqrt{h} + o_p(\sqrt{h}), \quad (49)$$

where $z_\alpha = \Phi^{-1}(\alpha)$. Consequently,

$$\beta_{\alpha,t}(h) \longrightarrow \frac{1}{2} \quad (h \downarrow 0) \quad (50)$$

whenever the elasticity is well defined.

Proof sketch. From (9),

$$\frac{X_{t+h} - X_t}{\sqrt{h}} = \frac{1}{\sqrt{h}} \int_t^{t+h} \sqrt{v_s} dW_s^S + o_p(1).$$

Local continuity of v implies that the martingale term differs from $\sqrt{v_t}(W_{t+h}^S - W_t^S)/\sqrt{h}$ by $o_p(1)$. The latter converges stably to $\sqrt{v_t}Z$ with $Z \sim N(0,1)$. Quantile convergence at a strictly positive normal density gives (48); uniform integrability of the locally Gaussian tail yields (49). $\square$

Proposition 3.5 is deliberately a leading-order statement. It does not imply that square-root transport is accurate at five, ten or forty days. In the mixed rough model, the future variance path changes by order h^H , return–volatility leverage creates conditional skewness, and randomness in integrated variance creates conditional excess kurtosis. Those effects enter at higher orders than the leading $\sqrt{h}$ term but can be material at operational holding periods. Section 4 derives the first rough corrections explicitly.

3.8 Physical versus risk-neutral measures in derivative risk

VaR and ES are defined under $\mathbb{P}$, whereas a derivative mark is a risk-neutral valuation. The distinction is particularly important in rough volatility because a model calibrated to the implied-volatility surface under $\mathbb{Q}$ embeds a volatility risk premium and should not be used mechanically as a historical P&L generator. Recent work on risk premia in rough-volatility models makes the distinction explicit (Bonesini et al., 2025).

Our implementation therefore separates the two tasks:

$\mathbb{P}$: state dynamics and loss-scenario probabilities,
 $\mathbb{Q}$: conditional derivative revaluation.

(51)

In a full-revaluation simulation, the path of (S, v) over the risk horizon is generated under the physical parameters. At the horizon, the derivative is repriced using a risk-neutral pricing map and a specified mapping from the physical volatility state to the risk-neutral state. A practical implementation can consider either a parsimonious constant volatility-risk premium or a richer state-dependent mapping. This separation allows the model to exploit option information without confusing risk-neutral calibration with physical tail forecasting.

3.9 Nested benchmark models and testable restrictions

The mixed specification nests several useful benchmarks:

1. If $\sigma_R = 0$, volatility is driven by a single diffusive OU log-volatility factor.
2. If $\sigma_D = 0$, the model reduces to a one-factor tempered rough log-volatility model.
3. If $H \uparrow 1/2$, K_R converges to an exponential kernel and the model approaches a two-factor Markov Gaussian log-volatility specification.
4. If $\rho_{SD} = \rho_{SR} = 0$, returns are conditionally Gaussian variance mixtures and all conditional asymmetry generated by leverage disappears.
5. If both $\sigma_D = \sigma_R = 0$, the model collapses locally to the constant-volatility Gaussian benchmark underlying exact square-root scaling.

These restrictions make the empirical analysis a model-discrimination exercise rather than a comparison of unrelated specifications. They also permit likelihood-ratio, information- criterion, forecasting and encompassing comparisons between rough, non-rough and mixed alternatives.

3.10 Identification map and economic interpretation

The model is deliberately parameterized so that different parts of the data are informative about different mechanisms. At a heuristic level,

$$\underbrace{H}_{\text{local slope}} \quad \underbrace{\sigma_R}_{\text{rough amplitude}} \quad \underbrace{\kappa_R}_{\text{rough persistence}} \quad \underbrace{\sigma_D, \kappa_D}_{\text{diffusive scale/persistence}} \quad \underbrace{\rho_{SR}, \rho_{SD}}_{\text{leverage}} \quad (52)$$

leave distinct signatures on the volatility variogram, autocovariance and return–volatility cross moments. Specifically, the short-lag slope of the log-volatility variogram is informative about $2H$, the curvature of the variogram across scales informs the rough–diffusive mixture, the long-lag decay is informative about mean reversion, and signed return–volatility cross moments inform leverage. This separation is imperfect in finite samples, especially when volatility is observed with measurement error. The estimation procedure in Section 5 therefore treats identification strength and measurement error as part of the empirical output rather than as secondary implementation details.

The main economic interpretation can now be stated succinctly. The model does not posit a new universal law of the form $ES(h) \propto h^H$. It posits that the *deviation* of finite-horizon tail risk from its local diffusive benchmark is generated by a volatility process whose dominant source of variation changes with the horizon. The next section shows how this mechanism produces explicit rough corrections to VaR and Expected Shortfall.

4 Tail-Risk Asymptotics and Horizon Scaling

Section 3 establishes that rough volatility does not replace the local $\sqrt{h}$ scaling of a continuous semimartingale return. The relevant question is of higher order: how does roughness deform the

conditional loss distribution around its Gaussian limit, and how do those deformations propagate into VaR, Expected Shortfall (ES) and horizon transport? A rigorous answer requires two logically distinct ingredients. First, one needs local cumulant expansions for the standardized loss. Second, one needs a regularity result strong enough to turn those cumulants into a *weighted, uniform* Edgeworth expansion; pointwise convergence of distribution functions is not sufficient for ES because ES integrates the entire upper tail.

This section makes that separation explicit. The technical structure follows the same principle as short-time density expansions for continuous stochastic-volatility models in El Euch, Fukasawa, et al. (2019): conditional Gaussian smoothing provides a smooth density, while a stochastic expansion of the volatility functional determines the Hermite corrections. More general Wiener-functional expansion results such as Tudor and Yoshida (2023) provide an alternative route. We impose uniformity only on a compact Hurst set bounded away from zero and one half; no claim is made that the expansion is uniform as $H \downarrow 0$ or $H \uparrow 1/2$.

4.1 Linearized loss, forward-variance scale, and localization

For a locally linear position with signed dollar exposure $\mathcal{E}_t \neq 0$, define the martingale component of the loss over $[t, t+h]$ by

$$M_{t,h} := -\mathcal{E}_t \int_t^{t+h} \sqrt{v_s} dW_s^S. \quad (53)$$

Its conditional mean is zero and, by Itô isometry,

$$a_{t,h}^2 := \mathbb{E}_t[M_{t,h}^2] = \mathcal{E}_t^2 V_{t,h}, \quad V_{t,h} := \int_0^h \xi_t(s) ds, \quad \xi_t(s) := \mathbb{E}_t[v_{t+s}]. \quad (54)$$

Thus $a_{t,h}$ is an exact conditional standard-deviation scale for the martingale loss, not a spot-volatility approximation. We standardize by

$$Z_{t,h} := \frac{M_{t,h}}{a_{t,h}}, \quad \mathbb{E}_t[Z_{t,h}] = 0, \quad \mathbb{E}_t[Z_{t,h}^2] = 1. \quad (55)$$

Using $a_{t,h}$ rather than $|\mathcal{E}_t|\sqrt{v_t h}$ absorbs the predictable movement of the conditional forward-variance curve, which can itself be of rough order.

The full locally linear loss can be written

$$L_{t,h} = m_{t,h} + M_{t,h} + R_{t,h}, \quad (56)$$

where $m_{t,h}$ is an $\mathcal{F}_t$ -measurable location term generated by drift and carry and $R_{t,h}$ collects higher-order price-linearization terms. Under the regularity conditions below, $m_{t,h} = O_p(h)$ and $R_{t,h} = O_p(h)$ for a cash position. Hence those terms are negligible relative to the first rough-leverage correction, whose absolute order is $h^{H+1/2}$, because $H < 1/2$. They need not be negligible relative to a pure kurtosis correction of order $h^{2H+1/2}$ when $H \geq 1/4$; this distinction will be stated explicitly rather than suppressed.

We use the following localization throughout. Let Θ be a compact parameter set with

$$0 < \underline{H} \leq H \leq \overline{H} < \frac{1}{2}, \quad 0 < \underline{\kappa} \leq \kappa_D, \kappa_R \leq \overline{\kappa}, \quad 0 \leq \sigma_D, \sigma_R \leq \overline{\sigma}, \quad (57)$$

and suppose the correlation matrix $\mathbf{R}$ in (7) satisfies

$$\inf_{\theta \in \Theta} \lambda_{\min}(\mathbf{R}) \geq r_* > 0. \quad (58)$$

In particular, the volatility block $\mathbf{R}_v$ is invertible. Write $\mathbf{R}_v$ and $\mathbf{r}$ as in (29) and define the residual return variance $q_S = 1 - \mathbf{r}'\mathbf{R}_v^{-1}\mathbf{r}$ as in (30).

Assumption 4.1 (Uniform nondegeneracy and local moments). *There are constants $r_* > 0$, $q_* > 0$, $h_* > 0$ and, for every $p > 1$, a finite C_p such that, uniformly over $\theta \in \Theta$ and $0 < h \leq h_*$,*

$$\lambda_{\min}(\mathbf{R}) \geq r_*, \quad (59)$$

$$q_S \geq q_*, \quad (60)$$

$$\left\| \frac{1}{h} \int_0^h v_{t+s} ds \right\|_{p,t} + \left\| \left\{ \frac{1}{h} \int_0^h q_S v_{t+s} ds \right\}^{-1} \right\|_{p,t} \leq C_p, \quad (61)$$

where $\|X\|_{p,t} := (\mathbb{E}_t|X|^p)^{1/p}$. The process $\sqrt{v}$ is predictable, $\mathbb{E}_t \int_t^{t+h} v_s ds < \infty$, the drift μ is locally bounded in conditional L^p , and the current variance is positive almost surely.

Assumption 4.1 is the two-factor analogue of the Gaussian smoothing condition used in short-time stochastic-volatility density expansions. The strict bound $q_S \geq q_*$ excludes the singular case in which the return Brownian motion is fully spanned by the volatility shocks. It is stronger than needed for a central limit theorem, but it is useful for a smooth density expansion and for tail integration.

Assumption 4.2 (Local kernel regularity). *For some $\eta > 0$, uniformly over $\theta \in \Theta$,*

$$K_R(u) = c_{H,\kappa_R} u^{H-1/2} \{1 + r_R(u)\}, \quad |r_R(u)| + u|r'_R(u)| \leq C u^\eta, \quad (62)$$

$$K_D(u) = K_D(0) + O(u), \quad K_D(0) = \sqrt{2\kappa_D}, \quad (63)$$

as $u \downarrow 0$. Both kernels are square-integrable on $(0, \infty)$ and have enough integrable derivatives away from zero to justify stochastic Fubini and differentiation under the integral sign in the local expansions below.

The exponentially tempered kernels in Section 3 satisfy Assumption 4.2. Define the local rough and diffusive amplitudes

$$\lambda_R := \sigma_R c_{H,\kappa_R}, \quad \lambda_D := \sigma_D K_D(0) = \sigma_D \sqrt{2\kappa_D}. \quad (64)$$

Assumption 4.3 (Local forward-variance regularity). *On each localization set $\{v_t \in [M^{-1}, M]\}$ and for every finite p , the conditional forward variance is locally continuous in the sense that*

$$\sup_{0 \leq s \leq h} \left| \frac{\xi_t(s)}{v_t} - 1 \right| = o_p(1), \quad (65)$$

and the conditional L^p stochastic Taylor expansion of $\sqrt{v_{t+s}/\xi_t(s)}$ obtained from the Gaussian–

Volterra exponential is uniform for $0 \leq s \leq h$. In particular, its first fresh rough term is

$$\frac{\lambda_R}{2} \int_0^s (s-u)^{H-1/2} dB_{t+u}^R, \quad (66)$$

with remainder $O_{L^p,t}(s^{2H}) + O_{L^p,t}(s^{1/2})$ uniformly over Θ .

For the exponentially tempered Gaussian kernels in Section 3, Assumption 4.3 is the local regularity condition needed to replace the conditional volatility scale by its current value in contraction calculations. It is stated explicitly because the coefficient calculations below use more than continuity in probability of the volatility path.

4.2 A uniform Edgeworth condition and why it is sufficient for ES

Let

$$\psi_{t,h}(u) := \mathbb{E}_t[e^{iuZ_{t,h}}] \quad (67)$$

be the conditional characteristic function of the centered, variance-one loss. The next assumption is deliberately stated at the characteristic-function level. It is strong enough to imply a weighted density expansion by Fourier inversion and weak enough to be verified from the stochastic expansion of Gaussian–Volterra functionals. In particular, it does not assume the VaR or ES expansion that we seek to prove.

Assumption 4.4 (Uniform Fourier–Edgeworth remainder). *Fix an integer $m \geq 3$. For a deterministic rate $r_h \downarrow 0$, suppose there is a polynomial $P_{t,h}(iu)$ of fixed degree, with $\mathcal{F}_t$ -measurable coefficients bounded in probability uniformly over Θ , such that*

$$\psi_{t,h}(u) = e^{-u^2/2} \{1 + P_{t,h}(iu)\} + \mathcal{R}_{t,h}(u), \quad (68)$$

with $P_{t,h}(0) = 0$, and, for $j = 0, \dots, m$,

$$\sup_{\theta \in \Theta} \int_{\mathbb{R}} \left| \partial_u^j \mathcal{R}_{t,h}(u) \right| du = o_p(r_h). \quad (69)$$

We additionally require $\mathcal{R}_{t,h} \in W^{m,1}(\mathbb{R})$ almost surely after localization, so that the boundary terms generated by Fourier integration by parts vanish. The same bounds hold on $\{v_t \in [M^{-1}, M]\}$ for every finite M , with constants independent of h .

Condition (69) is the technical device that makes the eventual ES expansion legitimate: it delivers polynomially weighted control of the density remainder, which is much stronger than convergence of a few moments. Such weighted density control is proved for general continuous stochastic-volatility models, including rough Bergomi, in El Euch, Fukasawa, et al. (2019). For general Wiener functionals, cumulant-based high-order density expansions can alternatively be justified by the framework of Tudor and Yoshida (2023).

Let H_k denote the probabilists' Hermite polynomials,

$$H_2(x) = x^2 - 1, \quad H_3(x) = x^3 - 3x, \quad H_4(x) = x^4 - 6x^2 + 3, \quad (70)$$

with $\{H_k(x)\phi(x)\}' = -H_{k+1}(x)\phi(x)$.

Theorem 4.5 (Weighted uniform Edgeworth inversion). *Under Assumption 4.4, $Z_{t,h}$ admits a conditional density $f_{t,h}$ and, if $q_{t,h}$ is the inverse Fourier transform of $e^{-u^2/2}\{1 + P_{t,h}(iu)\}$, then*

$$\sup_{\theta \in \Theta} \sup_{x \in \mathbb{R}} (1 + |x|)^m |f_{t,h}(x) - q_{t,h}(x)| = o_p(r_h). \quad (71)$$

Consequently, for the corresponding distribution functions $F_{t,h}$ and $Q_{t,h}$,

$$\sup_{\theta \in \Theta} \sup_{x \in \mathbb{R}} (1 + |x|)^{m-1} |F_{t,h}(x) - Q_{t,h}(x)| = o_p(r_h). \quad (72)$$

Proof. Fourier inversion gives

$$f_{t,h}(x) - q_{t,h}(x) = (2\pi)^{-1} \int_{\mathbb{R}} e^{-iux} \mathcal{R}_{t,h}(u) du.$$

Multiplication by x^j corresponds to j integrations by parts in u . Therefore (69) implies, for $j = 0, \dots, m$,

$$\sup_x |x|^j |f_{t,h}(x) - q_{t,h}(x)| = o_p(r_h).$$

Combining these finitely many bounds yields (71). Integrating the density remainder and using $m \geq 3$ gives (72). $\square$

The next theorem is the formal passage from a uniform distributional expansion to VaR and ES. It is stated uniformly over a compact set of confidence levels because the empirical analysis uses several fixed high-confidence levels. We do *not* claim uniformity as $\alpha \uparrow 1$; that would require a moderate- or large-deviation analysis rather than a central Edgeworth expansion.

Theorem 4.6 (Uniform transfer from Edgeworth tails to VaR and ES). *Let $\mathcal{A} = [\alpha_-, \alpha_+] \subseteq (1/2, 1)$ and $z_\alpha = \Phi^{-1}(\alpha)$. Suppose $F_{t,h}$ is conditionally continuous, $\mathbb{E}_t|Z_{t,h}| < \infty$, and, for some $\delta > 0$ and $\varepsilon_h \downarrow 0$,*

$$F_{t,h}(x) = \Phi(x) + \varepsilon_h G_t(x) + R_{t,h}(x). \quad (73)$$

Let K be a fixed neighborhood of $\{z_\alpha : \alpha \in \mathcal{A}\}$. Assume G_t is continuously differentiable on K and

$$\sup_{\theta \in \Theta} \left\{ \|G_t\|_{C^1(K)} + \int_{z_{\alpha_-}}^{\infty} |G_t(x)| dx \right\} = O_p(1), \quad (74)$$

and suppose

$$\sup_{\theta \in \Theta} \sup_{x \geq z_{\alpha_-} - 1} (1 + x)^{1+\delta} |R_{t,h}(x)| = o_p(\varepsilon_h). \quad (75)$$

Then, uniformly in $\alpha \in \mathcal{A}$,

$$q_{\alpha,t}(h) := F_{t,h}^{-1}(\alpha) = z_\alpha - \varepsilon_h \frac{G_t(z_\alpha)}{\phi(z_\alpha)} + o_p(\varepsilon_h), \quad (76)$$

$$e_{\alpha,t}(h) := \mathbb{E}_t[Z_{t,h} \mid Z_{t,h} \geq q_{\alpha,t}(h)] = c_\alpha - \frac{\varepsilon_h}{1 - \alpha} \int_{z_\alpha}^{\infty} G_t(x) dx + o_p(\varepsilon_h), \quad (77)$$

where

$$c_\alpha := \frac{\phi(z_\alpha)}{1 - \alpha}. \quad (78)$$

For any affine loss $Y_{t,h} = b_{t,h} + a_{t,h}Z_{t,h}$ with $a_{t,h} > 0$,

$$\text{VaR}_{\alpha,t}(Y) = b_{t,h} + a_{t,h}q_{\alpha,t}(h), \quad \text{ES}_{\alpha,t}(Y) = b_{t,h} + a_{t,h}e_{\alpha,t}(h). \quad (79)$$

Proof. Because $\mathcal{A}$ is compactly contained in $(0, 1)$, $\inf_{\alpha \in \mathcal{A}} \phi(z_\alpha) > 0$. Uniform inversion of (73) around z_α therefore gives (76). Because $F_{t,h}$ is continuous and has a finite first moment,

$$e_{\alpha,t} = q_{\alpha,t} + \frac{1}{1-\alpha} \int_{q_{\alpha,t}}^{\infty} \{1 - F_{t,h}(x)\} dx.$$

The first-order displacement of the lower integration limit cancels the first-order quantile displacement. The weighted bound (75) implies

$$\sup_{\alpha \in \mathcal{A}} \left| \int_{z_\alpha}^{\infty} R_{t,h}(x) dx \right| = o_p(\varepsilon_h),$$

which proves (77). Affine equivariance gives (79). $\square$

Remark 4.7 (Why moment matching alone is not enough). Knowing only that skewness is $O(h^H)$ or kurtosis is $O(h^{2H})$ does not, by itself, justify a Cornish–Fisher or ES expansion. Theorem 4.6 requires control of the distributional remainder over the tail that is integrated by ES. Assumption 4.4 and Theorem 4.5 provide that missing step.

4.3 Local cumulants generated by rough leverage

We next identify the first conditional cumulants of $Z_{t,h}$. Let

$$\varsigma_t := -\text{sgn}(\mathcal{E}_t), \quad (80)$$

so that $\varsigma_t = -1$ for a long cash position. Conditional on $\mathcal{F}_t$, the fresh rough innovation over $[t, t+s]$ satisfies, uniformly over Θ ,

$$\delta Y_{t+s}^R = \lambda_R \int_0^s (s-u)^{H-1/2} dB_{t+u}^R + o_{L^p,t}(s^H), \quad (81)$$

for every finite p . By Assumption 4.3, the first stochastic correction to $\sqrt{v_{t+s}}$ can be written uniformly in conditional L^p as

$$\sqrt{v_{t+s}} = \sqrt{\xi_t(s)} \left[1 + \frac{1}{2} \lambda_R \int_0^s (s-u)^{H-1/2} dB_{t+u}^R + O_{L^p,t}(s^{2H}) + O_{L^p,t}(s^{1/2}) \right]. \quad (82)$$

The $O(s^{1/2})$ term contains the diffusive volatility innovation; predictable movement of the Volterra factors is absorbed into $\xi_t(s)$ and hence into $a_{t,h}$.

Lemma 4.8 (Gaussian contraction identity). *Let*

$$I_h^R := \int_0^h \left\{ \int_0^s (s-u)^{H-1/2} dB_{t+u}^R \right\} dW_{t+s}^S.$$

Then

$$\mathbb{E}_t \left[(W_{t+h}^S - W_t^S)^2 I_h^R \right] = \frac{2\rho_{SR}}{(H + \frac{1}{2})(H + \frac{3}{2})} h^{H+3/2}. \quad (83)$$

Proof. After shifting time to zero, use $W_h^{S^2} - h = 2 \int_0^h W_s^S dW_s^S$. Itô isometry for correlated Brownian motions gives

$$\mathbb{E}[W_s^S \int_0^s (s-u)^{H-1/2} dB_u^R] = \rho_{SR} \int_0^s (s-u)^{H-1/2} du = \frac{\rho_{SR}}{H+1/2} s^{H+1/2}.$$

A final integration over $s \in [0, h]$ yields (83). $\square$

Theorem 4.9 (Leading rough-leverage skewness). *Under Assumptions 4.1–4.3, if $\lambda_R \rho_{SR} \neq 0$, the standardized martingale loss satisfies*

$$\kappa_{3,t}(h) := \mathbb{E}_t[Z_{t,h}^3] = \mathfrak{s}_{R,t} h^H + o_p(h^H), \quad (84)$$

where

$$\mathfrak{s}_{R,t} = -\text{sgn}(\mathcal{E}_t) \frac{3\lambda_R \rho_{SR}}{(H + \frac{1}{2})(H + \frac{3}{2})}. \quad (85)$$

For a long equity position with $\rho_{SR} < 0$, the coefficient is positive and the conditional loss distribution is right-skewed.

Proof. Insert (82) into (53). The leading term is Gaussian and the first rough correction is a second-order Wiener integral multiplied by $\lambda_R/2$. Expanding the third moment gives three identical contractions between two leading Gaussian terms and one rough correction. Lemma 4.8 therefore yields (85). After normalization by $a_{t,h}^3$, the remaining rough quadratic terms are of order h^{2H} , the diffusive-leverage term is of order $h^{1/2}$, and rough–diffusive cross terms are of order at least $h^{H+1/2}$. Since $0 < H < 1/2$, each is $o(h^H)$. $\square$

Remark 4.10 (Diffusive leverage as a nested benchmark). If the rough leverage channel is absent and the volatility specification is restricted so that the leading asymmetric contribution comes from the diffusive factor, then the same contraction argument gives

$$\kappa_{3,t}(h) = \mathfrak{s}_{D,t} h^{1/2} + o_p(h^{1/2}), \quad \mathfrak{s}_{D,t} = -\text{sgn}(\mathcal{E}_t) \frac{3\lambda_D \rho_{SD}}{2}. \quad (86)$$

We do not combine (84) and (86) into a universal two-term expansion. Higher-order rough terms can occur between h^H and $h^{1/2}$ when H is small, so the robust statement for the mixed model is the leading rough term in Theorem 4.9.

4.4 Variance-mixture kurtosis without return–volatility leverage

To isolate the symmetric tail mechanism, assume throughout this subsection that

$$\rho_{SR} = \rho_{SD} = 0. \quad (87)$$

Then W^S is independent of the volatility filtration and, conditional on the future variance path, $M_{t,h}$ is Gaussian with variance $\mathcal{E}_t^2 IV_{t,h}$, where $IV_{t,h} = \int_t^{t+h} v_s ds$.

A conditional L^2 Taylor expansion of the exponential volatility, justified by Assumptions 4.2–

4.3, gives the leading random part of integrated variance as

$$IV_{t,h} - \mathbb{E}_t[IV_{t,h}] = v_t \lambda_R \int_0^h \int_0^s (s-u)^{H-1/2} dB_{t+u}^R ds + o_{L^2,t}(h^{H+1}), \quad (88)$$

where replacing v_t by the locally equivalent forward-variance level changes only the remainder. Stochastic Fubini gives

$$\int_0^h \int_0^s (s-u)^{H-1/2} dB_{t+u}^R ds = \frac{1}{H + \frac{1}{2}} \int_0^h (h-u)^{H+1/2} dB_{t+u}^R. \quad (89)$$

Theorem 4.11 (Local excess kurtosis from rough integrated variance). *Under Assumptions 4.1–4.3 and (87),*

$$\kappa_{4,t}(h) := \mathbb{E}_t[Z_{t,h}^4] - 3 = \mathfrak{k}_R h^{2H} + o_p(h^{2H}), \quad (90)$$

where

$$\boxed{\mathfrak{k}_R = \frac{3\lambda_R^2}{(H + \frac{1}{2})^2(2H + 2)} > 0.} \quad (91)$$

Proof. Since $M_{t,h} \mid \{v_s : t \leq s \leq t+h\}$ is centered Gaussian,

$$\mathbb{E}_t[M_{t,h}^4] = 3\mathcal{E}_t^4 \mathbb{E}_t[IV_{t,h}^2].$$

Using $a_{t,h}^2 = \mathcal{E}_t^2 \mathbb{E}_t[IV_{t,h}]$,

$$\kappa_{4,t}(h) = 3 \frac{\text{Var}_t(IV_{t,h})}{\mathbb{E}_t[IV_{t,h}]^2}.$$

Equations (88)–(89) and Itô isometry yield

$$\frac{\text{Var}_t(IV_{t,h})}{\mathbb{E}_t[IV_{t,h}]^2} = \frac{\lambda_R^2}{(H + \frac{1}{2})^2(2H + 2)} h^{2H} + o_p(h^{2H}),$$

which proves the result. $\square$

4.5 Uniform Edgeworth expansions in the two economically relevant regimes

The cumulant theorems above determine the coefficients but do not by themselves establish a density expansion. We now state the precise regularity needed to close that gap.

Assumption 4.12 (Model-specific Fourier remainder). *Under the parameter and kernel conditions above, Assumption 4.4 holds uniformly on Θ in the following two cases:*

1. if $\lambda_R \rho_{SR} \neq 0$, take $r_h = h^H$ and

$$P_{t,h}(iu) = \frac{\mathfrak{s}_{R,t}}{6} h^H (iu)^3; \quad (92)$$

2. if $\rho_{SR} = \rho_{SD} = 0$, take $r_h = h^{2H}$ and

$$P_{t,h}(iu) = \frac{\mathfrak{k}_R}{24} h^{2H} (iu)^4. \quad (93)$$

Remark 4.13 (Verification of Assumption 4.12). The local stochastic expansion, the conditional L^p bounds and the residual Gaussian component $q_S \geq q_*$ provide the ingredients used in existing short-time density expansions for stochastic-volatility Wiener functionals; see El Euch, Fukasawa, et al. (2019) and Tudor and Yoshida (2023). The present paper therefore states the Fourier remainder as an explicit maintained regularity condition. This avoids claiming that a few cumulant calculations alone prove a uniform Edgeworth expansion for the full two-factor tempered Volterra model. The compact restriction $H \in [\underline{H}, \overline{H}]$ remains essential for the uniform statements below.

Combining Theorems 4.5–4.9 yields

$$f_{t,h}(x) = \phi(x) \left\{ 1 + \frac{\mathfrak{s}_{R,t}}{6} h^H H_3(x) \right\} + o_{w,p}(h^H), \quad (94)$$

where $o_{w,p}$ denotes the weighted uniform remainder in (71). Hence

$$F_{t,h}(x) = \Phi(x) - \frac{\mathfrak{s}_{R,t}}{6} h^H H_2(x) \phi(x) + o_{w,p}(h^H). \quad (95)$$

In the no-leverage regime,

$$f_{t,h}(x) = \phi(x) \left\{ 1 + \frac{\mathfrak{k}_R}{24} h^{2H} H_4(x) \right\} + o_{w,p}(h^{2H}), \quad (96)$$

$$F_{t,h}(x) = \Phi(x) - \frac{\mathfrak{k}_R}{24} h^{2H} H_3(x) \phi(x) + o_{w,p}(h^{2H}). \quad (97)$$

4.6 Stability of VaR and ES under a small dependent remainder

The local Edgeworth expansions are derived for the martingale component $M_{t,h}$. To transfer them to the full loss

$$L_{t,h} = m_{t,h} + M_{t,h} + R_{t,h},$$

one must control the probability and, for Expected Shortfall, the magnitude of rare remainder events. Mere convergence $R_{t,h} = o_p(a_{t,h} r_h)$ is not sufficient for a quantile expansion, and it is not sufficient for an Expected-Shortfall expansion either if the exceptional realizations of $R_{t,h}$ can be arbitrarily large. The following two results state sufficient local conditions used below.

Throughout this subsection, write

$$\mathbb{P}_t(\cdot) := \mathbb{P}(\cdot \mid \mathcal{F}_t), \quad X_{t,h} := \frac{M_{t,h}}{a_{t,h}}, \quad U_{t,h} := \frac{R_{t,h}}{a_{t,h}},$$

where $m_{t,h}$ and $a_{t,h} > 0$ are $\mathcal{F}_t$ -measurable.

Proposition 4.14 (Conditional VaR stability under a small remainder). *Let $\mathcal{A} = [\alpha_-, \alpha_+] \in (0, 1)$ and let $r_h \downarrow 0$. Assume that, conditionally on $\mathcal{F}_t$, $X_{t,h}$ admits a density $f_{t,h}$ on a neighborhood of $\{q_{\alpha,t,h} : \alpha \in \mathcal{A}\}$, where $q_{\alpha,t,h} := \text{VaR}_{\alpha,t}(X_{t,h})$, and that there exist constants $\eta > 0$ and $\underline{f} > 0$ such that*

$$\mathbb{P} \left[\inf_{\alpha \in \mathcal{A}} \inf_{|x - q_{\alpha,t,h}| \leq \eta} f_{t,h}(x) \geq \underline{f} \right] \longrightarrow 1. \quad (98)$$

Suppose further that there exists a deterministic sequence $\delta_h \downarrow 0$ such that

$$\pi_{t,h} := \mathbb{P}_t(|R_{t,h}| > a_{t,h}\delta_h r_h) = o_p(r_h). \quad (99)$$

No independence between $M_{t,h}$ and $R_{t,h}$ is required. Then, uniformly in $\alpha \in \mathcal{A}$,

$$\text{VaR}_{\alpha,t}(L_{t,h}) = m_{t,h} + \text{VaR}_{\alpha,t}(M_{t,h}) + o_p(a_{t,h}r_h). \quad (100)$$

More precisely, on an event whose probability tends to one,

$$\sup_{\alpha \in \mathcal{A}} \frac{|\text{VaR}_{\alpha,t}(L_{t,h}) - m_{t,h} - \text{VaR}_{\alpha,t}(M_{t,h})|}{a_{t,h}} \leq \delta_h r_h + \frac{2\pi_{t,h}}{\underline{f}}. \quad (101)$$

Proof. Conditional translation invariance and positive homogeneity reduce the problem to comparing the quantiles of $X_{t,h} + U_{t,h}$ and $X_{t,h}$. Put $b_h := \delta_h r_h$ and, for fixed $\alpha \in \mathcal{A}$, write $q = q_{\alpha,t,h}$. On the event in (98), set

$$s_h := b_h + \frac{2\pi_{t,h}}{\underline{f}}.$$

Because $b_h = o(r_h)$ and $\pi_{t,h} = o_p(r_h)$, one has $s_h = o_p(r_h)$ and hence $s_h < \eta$ with probability tending to one. The inclusions

$$\{X_{t,h} \leq q + s_h - b_h\} \cap \{|U_{t,h}| \leq b_h\} \subseteq \{X_{t,h} + U_{t,h} \leq q + s_h\}$$

and

$$\{X_{t,h} + U_{t,h} \leq q - s_h\} \cap \{|U_{t,h}| \leq b_h\} \subseteq \{X_{t,h} \leq q - s_h + b_h\}$$

give, respectively,

$$\begin{aligned} \mathbb{P}_t(X_{t,h} + U_{t,h} \leq q + s_h) &\geq F_{t,h}\left(q + \frac{2\pi_{t,h}}{\underline{f}}\right) - \pi_{t,h} \geq \alpha, \\ \mathbb{P}_t(X_{t,h} + U_{t,h} \leq q - s_h) &\leq F_{t,h}\left(q - \frac{2\pi_{t,h}}{\underline{f}}\right) + \pi_{t,h} \leq \alpha. \end{aligned}$$

The inequalities use the local lower bound on the density. Therefore the conditional α -quantile of $X_{t,h} + U_{t,h}$ lies in $[q - s_h, q + s_h]$. The constants are uniform over $\alpha \in \mathcal{A}$, which proves (101) and hence (100). $\square$

Proposition 4.15 (Conditional ES stability under an integrable remainder). *Let $\mathcal{A} = [\alpha_-, \alpha_+] \Subset (0, 1)$ and let $r_h \downarrow 0$. Assume that $M_{t,h}$ and $R_{t,h}$ are conditionally integrable. Suppose that there exists a deterministic sequence $\delta_h \downarrow 0$ such that, with the same threshold used in Proposition 4.14,*

$$\tau_{t,h} := \frac{1}{a_{t,h}r_h} \mathbb{E}_t\left[|R_{t,h}| \mathbf{1}_{\{|R_{t,h}| > a_{t,h}\delta_h r_h\}}\right] = o_p(1). \quad (102)$$

No independence between $M_{t,h}$ and $R_{t,h}$ is required, and no density assumption on $M_{t,h}$ is needed

for this ES result. Then, uniformly in $\alpha \in \mathcal{A}$,

$$\text{ES}_{\alpha,t}(L_{t,h}) = m_{t,h} + \text{ES}_{\alpha,t}(M_{t,h}) + o_p(a_{t,h}r_h). \quad (103)$$

More precisely,

$$\sup_{\alpha \in \mathcal{A}} \frac{|\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h} - \text{ES}_{\alpha,t}(M_{t,h})|}{a_{t,h}r_h} \leq \frac{\delta_h + \tau_{t,h}}{1 - \alpha_+}. \quad (104)$$

Proof. For an integrable loss Y , conditional Expected Shortfall admits the conditional Rockafellar–Uryasev representation

$$\text{ES}_{\alpha,t}(Y) = \text{ess inf}_{q \in L^0(\mathcal{F}_t)} \left\{ q + \frac{1}{1 - \alpha} \mathbb{E}_t[(Y - q)_+] \right\}. \quad (105)$$

Since the positive-part map is one-Lipschitz,

$$|(x + r - q)_+ - (x - q)_+| \leq |r|,$$

for all $x, r, q \in \mathbb{R}$. Applying (105) first with $Y = M_{t,h} + R_{t,h}$ and then with $Y = M_{t,h}$ gives the conditional Lipschitz bound

$$|\text{ES}_{\alpha,t}(M_{t,h} + R_{t,h}) - \text{ES}_{\alpha,t}(M_{t,h})| \leq \frac{1}{1 - \alpha} \mathbb{E}_t|R_{t,h}|. \quad (106)$$

Splitting the conditional first moment at the threshold $a_{t,h}\delta_h r_h$ yields

$$\begin{aligned} \mathbb{E}_t|R_{t,h}| &\leq a_{t,h}\delta_h r_h + \mathbb{E}_t \left[|R_{t,h}| \mathbf{1}_{\{|R_{t,h}| > a_{t,h}\delta_h r_h\}} \right] \\ &= a_{t,h}r_h(\delta_h + \tau_{t,h}). \end{aligned}$$

Because $\alpha \leq \alpha_+$, $(1 - \alpha)^{-1} \leq (1 - \alpha_+)^{-1}$, which proves (104). Finally, conditional translation invariance of ES adds the $\mathcal{F}_t$ -measurable location $m_{t,h}$, proving (103). $\square$

Corollary 4.16 (Common conditional moment condition for VaR and ES). *Suppose*

$$a_{t,h} \asymp_p h^{1/2}, \quad r_h = h^H, \quad 0 < H < \frac{1}{2},$$

and, for some $p \geq 1$,

$$\|R_{t,h}\|_{p,t} := (\mathbb{E}_t|R_{t,h}|^p)^{1/p} = O_p(h). \quad (107)$$

Then

$$\frac{\mathbb{E}_t|R_{t,h}|}{a_{t,h}h^H} = O_p(h^{1/2-H}) = o_p(1), \quad (108)$$

so Proposition 4.15 applies for every $p \geq 1$.

If, in addition,

$$p > \frac{H}{\frac{1}{2} - H}, \quad (109)$$

then one may choose

$$0 < \kappa < \frac{1}{2} - H - \frac{H}{p}, \quad \delta_h = h^\kappa,$$

and conditional Markov inequality yields

$$\mathbb{P}_t(|R_{t,h}| > a_{t,h}h^{H+\kappa}) = O_p(h^{p(\frac{1}{2}-H-\kappa)}) = o_p(h^H). \quad (110)$$

Hence, under the local density condition (98), the same L^p assumption delivers both VaR and ES stability at the rough correction scale.

Corollary 4.17 (Transfer of rough VaR and ES expansions to the full loss). *Suppose, uniformly for $\alpha \in \mathcal{A}$,*

$$\text{VaR}_{\alpha,t}(M_{t,h}) = a_{t,h} \left[z_\alpha + B_{\alpha,t}^V r_h + o_p(r_h) \right], \quad (111)$$

$$\text{ES}_{\alpha,t}(M_{t,h}) = a_{t,h} \left[c_\alpha + B_{\alpha,t}^E r_h + o_p(r_h) \right]. \quad (112)$$

If Propositions 4.14 and 4.15 hold, then

$$\text{VaR}_{\alpha,t}(L_{t,h}) = m_{t,h} + a_{t,h} \left[z_\alpha + B_{\alpha,t}^V r_h + o_p(r_h) \right], \quad (113)$$

$$\text{ES}_{\alpha,t}(L_{t,h}) = m_{t,h} + a_{t,h} \left[c_\alpha + B_{\alpha,t}^E r_h + o_p(r_h) \right]. \quad (114)$$

If moreover $m_{t,h} = o_p(a_{t,h}r_h)$, then the location term is negligible at the same correction order.

Proposition 4.18 (Stability and first-order expansion of normalized Expected Shortfall). *Let $\mathcal{A} = [\alpha_-, \alpha_+] \Subset (0, 1)$ and let $r_h \downarrow 0$. Consider*

$$L_{t,h} = m_{t,h} + M_{t,h} + R_{t,h},$$

where $m_{t,h}$ and $a_{t,h} > 0$ are $\mathcal{F}_t$ -measurable. Assume that $M_{t,h}$ and $R_{t,h}$ are conditionally integrable and that the martingale component satisfies, uniformly in $\alpha \in \mathcal{A}$,

$$\frac{\text{ES}_{\alpha,t}(M_{t,h})}{a_{t,h}} = c_\alpha + B_{\alpha,t}^E r_h + o_p(r_h), \quad c_\alpha := \frac{\phi(z_\alpha)}{1 - \alpha}, \quad (115)$$

where $z_\alpha = \Phi^{-1}(\alpha)$.

Suppose moreover that there exists a deterministic sequence $\delta_h \downarrow 0$ such that

$$\tau_{t,h} := \frac{1}{a_{t,h}r_h} \mathbb{E}_t \left[|R_{t,h}| \mathbf{1}_{\{|R_{t,h}| > a_{t,h}\delta_h r_h\}} \right] = o_p(1). \quad (116)$$

Then, uniformly in $\alpha \in \mathcal{A}$,

$$\boxed{\frac{\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h}}{a_{t,h}} = c_\alpha + B_{\alpha,t}^E r_h + o_p(r_h).} \quad (117)$$

Equivalently, after normalization by the matched-Gaussian Expected Shortfall,

$$\boxed{\frac{\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h}}{a_{t,h}c_\alpha} = 1 + A_{\alpha,t}^E r_h + o_p(r_h), \quad A_{\alpha,t}^E := \frac{B_{\alpha,t}^E}{c_\alpha}.} \quad (118)$$

More precisely,

$$\sup_{\alpha \in \mathcal{A}} \left| \frac{\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h}}{a_{t,h}} - \frac{\text{ES}_{\alpha,t}(M_{t,h})}{a_{t,h}} \right| \leq \frac{r_h}{1 - \alpha_+} (\delta_h + \tau_{t,h}). \quad (119)$$

Proof. Conditional translation invariance gives

$$\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h} = \text{ES}_{\alpha,t}(M_{t,h} + R_{t,h}).$$

The conditional L^1 -Lipschitz property of Expected Shortfall yields

$$|\text{ES}_{\alpha,t}(M_{t,h} + R_{t,h}) - \text{ES}_{\alpha,t}(M_{t,h})| \leq \frac{1}{1 - \alpha} \mathbb{E}_t |R_{t,h}|.$$

Splitting the first moment at the threshold $a_{t,h}\delta_h r_h$ gives

$$\mathbb{E}_t |R_{t,h}| \leq a_{t,h}\delta_h r_h + \mathbb{E}_t \left[|R_{t,h}| \mathbf{1}_{\{|R_{t,h}| > a_{t,h}\delta_h r_h\}} \right] = a_{t,h}r_h(\delta_h + \tau_{t,h}).$$

Since $\alpha \leq \alpha_+$, division by $a_{t,h}$ gives (119). Because $\delta_h + \tau_{t,h} = o_p(1)$, combining (119) with (115) proves (117). Finally, c_α is bounded away from zero on the compact set $\mathcal{A}$, so division by c_α yields (118). $\square$

Corollary 4.19 (Explicit rough-leverage correction). *Under the assumptions of Proposition 4.18 and the rough-leverage conditions of Theorem 4.22, take*

$$r_h = h^H, \quad 0 < H < \frac{1}{2}.$$

Then, uniformly in $\alpha \in \mathcal{A}$,

$$\boxed{\frac{\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h}}{a_{t,h}} = c_\alpha \left[1 + \frac{\mathfrak{s}_{R,t} z_\alpha}{6} h^H + o_p(h^H) \right]}, \quad (120)$$

where

$$\mathfrak{s}_{R,t} = -\text{sgn}(\mathcal{E}_t) \frac{3\chi_R}{(H + \frac{1}{2})(H + \frac{3}{2})}, \quad \chi_R := \lambda_R \rho_{SR}. \quad (121)$$

Hence

$$\boxed{\frac{\text{ES}_{\alpha,t}(L_{t,h}) - m_{t,h}}{a_{t,h}c_\alpha} = 1 - \text{sgn}(\mathcal{E}_t) \frac{\chi_R z_\alpha}{2(H + \frac{1}{2})(H + \frac{3}{2})} h^H + o_p(h^H)}. \quad (122)$$

Proof. For the martingale component, Theorem 4.22 gives

$$\text{ES}_{\alpha,t}(M_{t,h}) = a_{t,h} \left[c_\alpha + \frac{\mathfrak{s}_{R,t}}{6} \frac{z_\alpha \phi(z_\alpha)}{1 - \alpha} h^H + o_p(h^H) \right].$$

Since

$$c_\alpha = \frac{\phi(z_\alpha)}{1 - \alpha},$$

the correction coefficient equals $c_\alpha \mathfrak{s}_{R,t} z_\alpha / 6$. Proposition 4.18 transfers the expansion to the full centered loss. Substituting (121) gives (122). $\square$

Corollary 4.20 (Negligible location term). *Under the assumptions of Corollary 4.19, if*

$$\frac{m_{t,h}}{a_{t,h}} = o_p(h^H), \quad (123)$$

then the same expansion holds without centering:

$$\boxed{\frac{\text{ES}_{\alpha,t}(L_{t,h})}{a_{t,h}c_\alpha} = 1 - \text{sgn}(\mathcal{E}_t) \frac{\chi_R z_\alpha}{2(H + \frac{1}{2})(H + \frac{3}{2})} h^H + o_p(h^H)}. \quad (124)$$

In particular, if $m_{t,h} = O_p(h)$ and $a_{t,h} \asymp_p h^{1/2}$, then (123) follows automatically for every $H < 1/2$ because $h^{1/2} = o(h^H)$.

Remark 4.21 (Interpretation). Equation (122) separates the Gaussian benchmark from the rough tail correction. The constant

$$c_\alpha = \frac{\phi(z_\alpha)}{1 - \alpha}$$

is the Expected Shortfall of a standard normal loss beyond its α -quantile, whereas

$$- \text{sgn}(\mathcal{E}_t) \frac{\chi_R z_\alpha}{2(H + \frac{1}{2})(H + \frac{3}{2})}$$

is the structural first-order rough-leverage loading. For a long equity position, $\text{sgn}(\mathcal{E}_t) = 1$; therefore a negative return–volatility leverage product $\chi_R < 0$ generates a positive first-order ES add-on.

4.7 VaR and ES theorems with uniform confidence-level control

Theorem 4.22 (Rough-leverage correction to VaR and ES). *Let $\mathcal{A} = [\alpha_-, \alpha_+] \Subset (1/2, 1)$. Under Assumptions 4.1, 4.2 and 4.12, suppose $\lambda_R \rho_{SR} \neq 0$. Then, uniformly in $\alpha \in \mathcal{A}$ and $\theta \in \Theta$,*

$$\text{VaR}_{\alpha,t}(M_{t,h}) = a_{t,h} \left[z_\alpha + \frac{\mathfrak{s}_{R,t}}{6} (z_\alpha^2 - 1) h^H + o_p(h^H) \right], \quad (125)$$

$$\text{ES}_{\alpha,t}(M_{t,h}) = a_{t,h} \left[c_\alpha + \frac{\mathfrak{s}_{R,t}}{6} \frac{z_\alpha \phi(z_\alpha)}{1 - \alpha} h^H + o_p(h^H) \right]. \quad (126)$$

Proof. Equation (95) has the form (73) with $\varepsilon_h = h^H$ and $G_t(x) = -(\mathfrak{s}_{R,t}/6)H_2(x)\phi(x)$. Theorem 4.6 gives

$$-G_t(z)/\phi(z) = \frac{\mathfrak{s}_{R,t}}{6} H_2(z),$$

and, using $\int_z^\infty H_2(x)\phi(x)dx = z\phi(z)$,

$$- \int_z^\infty G_t(x)dx = \frac{\mathfrak{s}_{R,t}}{6} z\phi(z).$$

Affine scaling by $a_{t,h}$ proves (125)–(126). Transfer to the full loss $L_{t,h} = m_{t,h} + M_{t,h} + R_{t,h}$ follows from Corollary 4.17 under its explicit remainder conditions. $\square$

The preceding theorem is intentionally normalized by the exact forward-variance scale $a_{t,h}$. If one insists on the spot-volatility square-root benchmark, the change in the forward-variance curve contributes at the same order and must be retained.

Corollary 4.23 (Optional spot-volatility tangent representation). *Suppose, in addition to the assumptions of Theorem 4.22, that the forward-variance curve admits the extra local tangent condition: for an $\mathcal{F}_t$ -measurable random coefficient ϑ_t ,*

$$\frac{V_{t,h}}{v_t h} = 1 + \vartheta_t h^H + o_p(h^H). \quad (127)$$

Let $s_t = |\mathcal{E}_t| \sqrt{v_t}$. Then

$$\text{ES}_{\alpha,t}(M_{t,h}) = s_t \sqrt{h} \left[c_\alpha + \left\{ \frac{1}{2} c_\alpha \vartheta_t + \frac{\mathfrak{s}_{R,t}}{6} \frac{z_\alpha \phi(z_\alpha)}{1 - \alpha} \right\} h^H + o_p(h^H) \right]. \quad (128)$$

Condition (127) is an additional tangent assumption and is not implied by roughness alone. When it holds, the first deviation from spot square-root scaling contains a scale effect through the forward-variance curve and a distinct shape effect through rough leverage.

Theorem 4.24 (Pure rough variance-mixture correction). *Let $\mathcal{A} = [\alpha_-, \alpha_+] \Subset (1/2, 1)$. Under Assumptions 4.1, 4.2 and 4.12, suppose $\rho_{SR} = \rho_{SD} = 0$. Then, uniformly in $\alpha \in \mathcal{A}$ and $\theta \in \Theta$,*

$$\text{VaR}_{\alpha,t}(M_{t,h}) = a_{t,h} \left[z_\alpha + \frac{\mathfrak{k}_R}{24} (z_\alpha^3 - 3z_\alpha) h^{2H} + o_p(h^{2H}) \right], \quad (129)$$

$$\text{ES}_{\alpha,t}(M_{t,h}) = a_{t,h} \left[c_\alpha + \frac{\mathfrak{k}_R}{24} \frac{\phi(z_\alpha)(z_\alpha^2 - 1)}{1 - \alpha} h^{2H} + o_p(h^{2H}) \right]. \quad (130)$$

At conventional market-risk confidence levels $z_\alpha^2 > 1$, the shape correction to ES is positive.

Proof. Use (97) in Theorem 4.6. Since $\int_z^\infty H_3(x) \phi(x) dx = H_2(z) \phi(z)$, the stated coefficients follow. $\square$

Remark 4.25 (When is the kurtosis term the first correction to the *uncentered* risk?). Theorem 4.24 is a shape expansion around a matched Gaussian with the same conditional scale $a_{t,h}$. For the full uncentered loss, deterministic carry and drift are of relative order $h^{1/2}$. Therefore the rough kurtosis term h^{2H} is the first correction to the full VaR/ES only when $H < 1/4$ (or when the position is explicitly centered). If $H \geq 1/4$, drift, scale and diffusive corrections must be retained before or alongside the rough kurtosis term. This ordering issue is important in empirical implementations and is one reason we compare direct risk with a matched Gaussian benchmark rather than interpreting every deviation as “roughness.”

4.8 Estimable local risk statistics and a minimal identification result

The preceding theorems are structural, but the leading rough corrections depend on a much smaller set of parameter combinations than the full mixed-rough model. This observation is useful both econometrically and operationally. In the leverage regime, the first-order shape correction depends on H and on the product

$$\chi_R := \lambda_R \rho_{SR}, \quad (131)$$

rather than on λ_R and ρ_{SR} separately. In the symmetric no-leverage regime, the leading shape correction can be expressed through the conditional coefficient of variation of integrated variance.

Thus tail-risk functionals can remain well identified even when the structural decomposition of volatility is less precise.

Define the standardized return innovation

$$U_{t,h}^S := \frac{W_{t+h}^S - W_t^S}{\sqrt{h}} \quad (132)$$

and the fresh rough log-volatility innovation

$$U_{t,h}^R := \lambda_R \int_0^h (h-u)^{H-1/2} dB_{t+u}^R. \quad (133)$$

Its conditional variance and its cross moment with the standardized return innovation are, to first order,

$$\mathcal{S}_R(h) := \text{Var}_t(U_{t,h}^R) = \frac{\lambda_R^2}{2H} h^{2H}, \quad (134)$$

$$\mathcal{L}_R(h) := \text{Cov}_t(U_{t,h}^S, U_{t,h}^R) = \frac{\lambda_R \rho_{SR}}{H + \frac{1}{2}} h^H. \quad (135)$$

The quantity $\mathcal{L}_R(h)$ is a local leverage statistic. It can be estimated from returns and filtered innovations of log realized variance without separately identifying λ_R and ρ_{SR} .

Corollary 4.26 (Leverage-to-ES mapping). *Under the assumptions of Theorem 4.22,*

$$\kappa_{3,t}(h) = -\text{sgn}(\mathcal{E}_t) \frac{3}{H + \frac{3}{2}} \mathcal{L}_R(h) + o_p(h^H). \quad (136)$$

Consequently, uniformly for $\alpha \in \mathcal{A}$,

$$\boxed{\frac{\text{ES}_{\alpha,t}(M_{t,h})}{a_{t,h}c_\alpha} - 1 = -\text{sgn}(\mathcal{E}_t) \frac{z_\alpha}{2(H + \frac{3}{2})} \mathcal{L}_R(h) + o_p(h^H).} \quad (137)$$

For a long equity position, a negative return–volatility cross moment therefore maps directly into a positive ES shape add-on.

Proof. Substituting (135) into (85) gives (136). Divide (126) by $a_{t,h}c_\alpha$ and use $c_\alpha = \phi(z_\alpha)/(1-\alpha)$.

□

The no-leverage channel admits an equally transparent representation. Define

$$\mathcal{U}_{IV,t}(h) := \frac{\text{Var}_t(IV_{t,h})}{\mathbb{E}_t[IV_{t,h}]^2}. \quad (138)$$

Under return–volatility independence, Theorem 4.11 in fact uses the exact identity $\kappa_{4,t}(h) = 3\mathcal{U}_{IV,t}(h)$ before applying the local expansion.

Corollary 4.27 (Integrated-variance uncertainty and ES). *Under the assumptions of Theorem 4.24,*

$$\boxed{\frac{\text{ES}_{\alpha,t}(M_{t,h})}{a_{t,h}c_\alpha} - 1 = \frac{z_\alpha^2 - 1}{8} \mathcal{U}_{IV,t}(h) + o_p(h^{2H}).} \quad (139)$$

Hence the symmetric rough correction can be estimated from uncertainty about integrated variance itself, without first translating that uncertainty into an abstract kurtosis parameter.

These representations suggest a useful separation between *structural identification* and *risk-functional identification*. Full estimation of $(H, \lambda_R, \rho_{SR}, \kappa_R, \dots)$ is required to simulate the complete mixed-rough model, but the leading ES correction may be estimated more robustly from $(H, \mathcal{L}_R(h))$ or from $\mathcal{U}_{IV,t}(h)$.

Proposition 4.28 (Minimal identification under the local power-law model). *Let $0 < h_1 < h_2$ and suppose that, over the pilot window, the leading local moment relations*

$$\mathcal{S}_R(h) = \frac{\lambda_R^2}{2H} h^{2H}, \quad \mathcal{L}_R(h) = \frac{\lambda_R \rho_{SR}}{H + \frac{1}{2}} h^H \quad (140)$$

are imposed as the local approximation. If $\mathcal{S}_R(h_j) > 0$, $j = 1, 2$, then $(H, \lambda_R, \rho_{SR})$ is identified by

$$H = \frac{1}{2} \frac{\log\{\mathcal{S}_R(h_2)/\mathcal{S}_R(h_1)\}}{\log(h_2/h_1)}, \quad (141)$$

$$\lambda_R = \left\{ 2H \mathcal{S}_R(h_1) h_1^{-2H} \right\}^{1/2}, \quad (142)$$

$$\rho_{SR} = \frac{(H + \frac{1}{2}) \mathcal{L}_R(h_1)}{\lambda_R h_1^H}. \quad (143)$$

For the full tempered model at fixed nonzero horizons these formulas define local effective parameters rather than exact structural parameters. Their structural consistency requires a shrinking-horizon argument, given in Section 5.3.

Proof. The ratio of (134) at h_2 and h_1 cancels $\lambda_R^2/(2H)$, which yields (141). Substitution back into (134) gives (142); inserting the resulting pair into (135) gives (143). $\square$

The closed-form identities in Proposition 4.28 are used as pilot estimators in Section 5. Section 7 then evaluates their finite-sample behavior and the resulting ES correction.

4.9 Horizon elasticity and a decomposition of the Rough Horizon Multiplier

The forward-variance normalization gives a useful exact decomposition of horizon scaling. For the martingale loss, define the *scale horizon multiplier*

$$\text{SHM}_t(h \mid h_0) := \frac{a_{t,h}}{a_{t,h_0} \sqrt{h/h_0}} = \sqrt{\frac{V_{t,h}/h}{V_{t,h_0}/h_0}}. \quad (144)$$

Consider a joint local-horizon regime $h_0 \downarrow 0$ with h/h_0 bounded away from zero and infinity. If the standardized ES has the generic expansion

$$e_{\alpha,t}(h) = c_\alpha \{1 + A_{\alpha,t} h^\delta + o_p(h^\delta)\}, \quad \delta > 0, \quad (145)$$

then, neglecting an $O(h)$ location term at the relevant order,

$$\text{RHM}_{\alpha,t}(h \mid h_0) = \text{SHM}_t(h \mid h_0) \frac{1 + A_{\alpha,t} h^\delta + o_p(h^\delta)}{1 + A_{\alpha,t} h_0^\delta + o_p(h_0^\delta)}. \quad (146)$$

Consequently, when both distortions are small,

$$\boxed{\text{RHM}_{\alpha,t}(h \mid h_0) - 1 = \{\text{SHM}_t(h \mid h_0) - 1\} + A_{\alpha,t}(h^\delta - h_0^\delta) + o_p(\Delta_{h,h_0})}, \quad (147)$$

where Δ_{h,h_0} collects the magnitudes of the scale and shape corrections. Equation (147) separates a *variance-term-structure effect* from a *non-Gaussian shape effect*. This decomposition is central for the empirical interpretation of the Rough Horizon Multiplier.

Under rough leverage,

$$\delta = H, \quad A_{\alpha,t}^{(L)} = \frac{\mathfrak{s}_{R,t} z_\alpha}{6}, \quad (148)$$

whereas under the no-leverage rough variance-mixture channel,

$$\delta = 2H, \quad A_{\alpha,t}^{(K)} = \frac{\mathfrak{k}_R}{24}(z_\alpha^2 - 1). \quad (149)$$

For the standardized shape component alone, the local elasticity is

$$\frac{d \log e_{\alpha,t}(h)}{d \log h} = \delta A_{\alpha,t} h^\delta + o_p(h^\delta). \quad (150)$$

If the variance scale is locally flat so that $a_{t,h} \propto \sqrt{h}$, the full ES elasticity is

$$\beta_{\alpha,t}(h) = \frac{1}{2} + \delta A_{\alpha,t} h^\delta + o_p(h^\delta). \quad (151)$$

In general, the full ES elasticity also contains the derivative of the forward-variance scale in (144). This distinction prevents the horizon term structure from being misread as a pure Hurst-exponent effect.

4.10 A structural decomposition of horizon-scaling risk

For practitioner interpretation, we retain three nested risk benchmarks. Let

$$ES_{\alpha,t}^G(h) := m_{t,h} + a_{t,h} c_\alpha, \quad (152)$$

$$ES_{\alpha,t}^{SV}(h) := \text{ES under a non-rough stochastic-volatility benchmark}, \quad (153)$$

$$ES_{\alpha,t}^{MR}(h) := \text{ES under the full mixed-rough model}. \quad (154)$$

The first benchmark is Gaussian but matched to the same conditional mean and forward-variance scale. It therefore isolates *shape* effects rather than confounding them with time variation in conditional variance.

We define the roughness-related add-on

$$\text{RRA}_{\alpha,t}(h) = ES_{\alpha,t}^{MR}(h) - ES_{\alpha,t}^{SV}(h) \quad (155)$$

and the horizon-transport add-on

$$\text{HTA}_{\alpha,t}(h \mid h_0) = ES_{\alpha,t}^{MR}(h) - ES_{\alpha,t}^{MR}(h_0) \sqrt{h/h_0}. \quad (156)$$

RRA asks how much risk is added by the rough/multi-scale volatility mechanism relative to a non-rough stochastic-volatility model. HTA asks how much risk is lost or added by transporting the *same* model's ES with a square-root rule instead of simulating the target horizon directly.

4.11 Nonlinear books and Greek transmission channels

The preceding theorems are local results for a linear exposure. A nonlinear portfolio must not be treated by simply replacing $a_{t,h}$ with a portfolio volatility. The P&L map itself changes the tail geometry. Using (38), write

$$L_{t,h}^G = L_\Delta + L_\Gamma + L_V + L_{Va} + L_{Vo} + L_\Theta, \quad (157)$$

where the terms denote Delta, Gamma, Vega, Vanna, Volga and Theta contributions. A natural position-level ES allocation is Euler allocation when the loss is positively homogeneous in position sizes. If w_j is the size of position j and differentiability holds,

$$ES_{\alpha,t}(h) = \sum_j w_j \frac{\partial ES_{\alpha,t}(h)}{\partial w_j}. \quad (158)$$

Because Greek channels interact, an exact additive decomposition by Greek need not exist. The empirical implementation therefore complements Euler position allocation with Shapley-style or sequential-increment channel attribution.

The testable implication is not that roughness must increase every nonlinear portfolio's ES. Gamma can magnify or cushion losses depending on portfolio sign; Vanna can amplify or offset return–volatility leverage; and Volga can become dominant in stress. The relevant hypothesis is that substantial volatility and curvature exposure makes the difference between direct mixed-rough ES and simpler horizon-scaled approximations economically larger.

4.12 Implications for estimation and risk implementation

The theory yields four implications that guide the empirical analysis. First, negative rough-factor leverage for a long equity position generates a positive ES shape correction of relative order h^H . Second, when return–volatility leverage is absent, uncertainty about integrated variance produces a positive high-confidence ES correction of order h^{2H} . Third, horizon distortion should be separated into a forward-variance scale component and a non-Gaussian shape component; the resulting tail-risk crossover need not coincide exactly with the volatility crossover h^* . Fourth, nonlinear books with material volatility and curvature exposure should exhibit larger differences between direct simulation and simple horizon transport than otherwise comparable linear portfolios.

Section 5 develops estimators for the structural and risk-functional parameters entering these implications.

5 Estimation under the Physical Measure

The empirical implementation must identify short-scale roughness, medium-horizon persistence, and return–volatility leverage without confusing these mechanisms with measurement error in realized volatility. This section therefore uses a staged physical-measure estimation strategy. The first stage estimates the volatility geometry from realized measures; the second stage estimates the joint return–volatility dependence; the third stage refines the parameters by matching direct-horizon return and tail moments. The staged approach is not required by the model, but it makes identification transparent and supplies robust starting values for the final joint criterion.

5.1 Observed returns and noisy volatility proxies

Let r_t denote the close-to-close log return and let RV_t be a realized-variance proxy constructed from intraday returns. We use

$$y_t = \log RV_t \quad (159)$$

as the primary volatility observable. Because RV_t estimates integrated rather than spot variance and because logarithms introduce Jensen bias, we employ the measurement equation

$$y_t = b_t + \log IV_t + e_t, \quad IV_t = \int_{t-1}^t v_s ds, \quad \mathbb{E}[e_t | \mathcal{G}_{t-1}] = 0, \quad (160)$$

where b_t is an analytic, estimated or absorbed log-measurement bias term and e_t is the remaining measurement disturbance. The baseline allows weak serial dependence in e_t and requires orthogonality to future Brownian innovations after filtering. This formulation is more conservative than identifying y_t directly with $\log v_t$ and addresses the debate on spurious roughness induced by noisy volatility measurement (Bolko et al., 2023; Cont and Das, 2024).

When several realized measures are available, the baseline uses a noise-robust realized kernel or pre-averaged measure and treats the standard realized variance as a robustness check. The empirical conclusions are required to survive changes in sampling frequency and volatility estimator.

5.2 Stage 0: closed-form pilot estimates for the rough tail-risk channel

The pilot stage is constructed from observable filtered moments. At fixed nonzero horizons, their probability limits are interpreted as *local effective targets*. Equality with the fresh Brownian–Volterra moments in Section 4 requires an additional measurement-link condition stated below. This distinction keeps filtering and de-noising assumptions visible rather than embedding them silently in the asymptotic theorem.

Before solving the full nonlinear GMM problem, we compute transparent pilot estimates of the parameters that enter the leading ES correction. This stage is intentionally simple: it is used for diagnostics, starting values and economic interpretation, not as the final estimator.

Let

$$\hat{u}_t(h) = y_{t+h} - \hat{\mathbb{E}}_t[y_{t+h}] \quad (161)$$

denote an h -horizon innovation from a preliminary state-space or linear volatility filter. After

subtracting the estimated measurement-error contribution and the predictable diffusive-factor contribution, define the rough-innovation variance estimate

$$\widehat{\mathcal{S}}_R(h) = \widehat{\text{Var}}\{\widehat{u}_t(h)\}_{\text{rough}}. \quad (162)$$

For two short horizons $h_1 < h_2$, the pilot Hurst estimate is

$$\widehat{H}^{\text{pilot}} = \frac{1}{2} \frac{\log\{\widehat{\mathcal{S}}_R(h_2)/\widehat{\mathcal{S}}_R(h_1)\}}{\log(h_2/h_1)}. \quad (163)$$

The local rough amplitude follows from

$$\widehat{\lambda}_R^{\text{pilot}} = \left\{ 2\widehat{H}^{\text{pilot}} \widehat{\mathcal{S}}_R(h_1) h_1^{-2\widehat{H}^{\text{pilot}}} \right\}^{1/2}. \quad (164)$$

To estimate the leverage channel, standardize the return by a current or forward-variance estimate,

$$\widehat{z}_t(h) = \frac{r_{t,t+h} - \widehat{\mu}_t h}{\sqrt{\widehat{V}_{t,h}}}, \quad (165)$$

and compute

$$\widehat{\mathcal{L}}_R(h) = \widehat{\text{Cov}}(\widehat{z}_t(h), \widehat{u}_t(h))_{\text{rough}}. \quad (166)$$

The pilot leverage estimate is then

$$\widehat{\rho}_{SR}^{\text{pilot}} = \frac{(\widehat{H}^{\text{pilot}} + \frac{1}{2})\widehat{\mathcal{L}}_R(h_1)}{\widehat{\lambda}_R^{\text{pilot}} h_1^{\widehat{H}^{\text{pilot}}}}. \quad (167)$$

A particularly robust object is the product

$$\widehat{\chi}_R = \widehat{\lambda}_R^{\text{pilot}} \widehat{\rho}_{SR}^{\text{pilot}}, \quad (168)$$

because this product, rather than its two factors separately, governs the first-order rough-leverage correction to ES.

The pilot estimates are accompanied by a moving-block bootstrap or stationary bootstrap. This is useful for two reasons. First, the same data are used to construct volatility proxies and return–volatility cross moments, so iid standard errors are inappropriate. Second, uncertainty in $\widehat{H}$ enters nonlinearly through $h^{\widehat{H}}$ and can be economically material at 20–120 day horizons. The bootstrap therefore propagates parameter uncertainty directly into the rough ES add-on.

The closed-form stage also provides a specification diagnostic. If estimates of H based on different pairs (h_1, h_2) vary sharply, or if the implied $\widehat{\rho}_{SR}$ repeatedly lies outside $[-1, 1]$, the one-component local approximation is inadequate. In that case the full mixed model, measurement error, or a different crossover location must be taken seriously before tail-risk conclusions are drawn.

Assumption 5.1 (Measurement link for structural interpretation). *Let $\widehat{m}_n$ denote the pilot moments produced by the chosen filtering and de-noising procedure and let their probability limit*

be $m_0 = (S_{1,0}, S_{2,0}, L_{1,0})'$. The fixed-horizon asymptotic theory below requires only existence of this limit. Whenever the closed-form structural interpretation is invoked, we additionally assume

$$(S_{1,0}, S_{2,0}, L_{1,0}) = (\mathcal{S}_R(h_1), \mathcal{S}_R(h_2), \mathcal{L}_R(h_1)). \quad (169)$$

Without (169), $(H^\dagger, \chi_R^\dagger)$ below remain local effective parameters of the observable pilot moments rather than primitive structural parameters.

5.3 Asymptotic theory for the pilot estimators

The closed-form pilot estimates are useful only if their sampling uncertainty is made explicit. This subsection provides two complementary asymptotic results. The first is a long-span result at fixed pilot horizons and is the one used for empirical standard errors. The second is a shrinking-horizon result that clarifies when the local pilot targets converge to the structural rough parameters and reveals an important rate loss for the structural rough-leverage coefficient.

For fixed $0 < h_1 < h_2$, define

$$m_0 = \begin{pmatrix} S_{1,0} \\ S_{2,0} \\ L_{1,0} \end{pmatrix} := \text{plim}_{n \rightarrow \infty} \hat{m}_n, \quad d_h := \log(h_2/h_1), \quad \ell_1 := \log h_1. \quad (170)$$

Let

$$\hat{m}_n = \begin{pmatrix} \hat{S}_{1,n} \\ \hat{S}_{2,n} \\ \hat{L}_{1,n} \end{pmatrix} \quad (171)$$

denote the corresponding noise-corrected sample moments after the preliminary volatility filter.

Assumption 5.2 (High-level asymptotic linearity of the pilot moments). *The observable filtered-moment estimator admits the high-level representation*

$$\sqrt{n}(\hat{m}_n - m_0) = \frac{1}{\sqrt{n}} \sum_{t=1}^n \psi_t + o_p(1), \quad (172)$$

where $\{\psi_t\}$ is a strictly stationary mean-zero vector process such that

$$\Omega_0 := \sum_{k=-\infty}^{\infty} \text{Cov}(\psi_0, \psi_k) \quad (173)$$

exists and is finite and positive definite, and

$$\frac{1}{\sqrt{n}} \sum_{t=1}^n \psi_t \Rightarrow N(0, \Omega_0). \quad (174)$$

The first-stage filter, measurement-error correction and standardization error are either $o_p(n^{-1/2})$ uniformly in the selected fixed horizons or are incorporated in the influence function ψ_t . In addition, $\mathbb{P}(\hat{S}_{1,n} > 0, \hat{S}_{2,n} > 0) \rightarrow 1$. This is a high-level assumption; it is not claimed to follow automatically from rough-volatility dynamics or from a generic mixing condition.

Assumption 5.2 is deliberately stated in influence-function form. It therefore covers standard mixing-based central limit theorems, HAC-corrected sample moments and generated volatility innovations, provided that the first-stage estimation error is included in ψ_t .

Define the fixed-horizon local targets

$$H^\dagger := g_H(m_0) = \frac{1}{2d_h} \log \left(\frac{S_{2,0}}{S_{1,0}} \right), \quad (175)$$

$$\chi_R^\dagger := g_\chi(m_0) = \left(H^\dagger + \frac{1}{2} \right) L_{1,0} h_1^{-H^\dagger}. \quad (176)$$

Under Assumption 5.1, when the local power-law moment restrictions

$$\mathcal{S}_R(h) = \frac{\lambda_R^2}{2H} h^{2H}, \quad \mathcal{L}_R(h) = \frac{\chi_R}{H + \frac{1}{2}} h^H$$

hold exactly at h_1 and h_2 , then $(H^\dagger, \chi_R^\dagger) = (H, \chi_R)$. Under a tempered or otherwise non-power-law rough kernel, $(H^\dagger, \chi_R^\dagger)$ should instead be interpreted as the local effective parameters associated with the selected pilot horizons.

Theorem 5.3 (Consistency and joint central limit theorem for $\hat{H}$ and $\hat{\chi}_R$). *Suppose Assumption 5.2 holds, $S_{1,0} > 0$, $S_{2,0} > 0$, and $H^\dagger$ lies in a compact subset of $(0, 1/2)$. Define*

$$\hat{H}_n = \frac{1}{2d_h} \log \left(\frac{\hat{S}_{2,n}}{\hat{S}_{1,n}} \right), \quad (177)$$

$$\hat{\chi}_{R,n} = \left(\hat{H}_n + \frac{1}{2} \right) \hat{L}_{1,n} h_1^{-\hat{H}_n}. \quad (178)$$

Then

$$(\hat{H}_n, \hat{\chi}_{R,n}) \xrightarrow{p} (H^\dagger, \chi_R^\dagger), \quad (179)$$

and

$$\sqrt{n} \begin{pmatrix} \hat{H}_n - H^\dagger \\ \hat{\chi}_{R,n} - \chi_R^\dagger \end{pmatrix} \implies N(0, V_{H\chi}), \quad V_{H\chi} = J_0 \Omega_0 J_0'. \quad (180)$$

Let $e_3 = (0, 0, 1)'$ and define

$$a_H = \begin{pmatrix} -\frac{1}{2d_h S_{1,0}} \\ \frac{1}{2d_h S_{2,0}} \\ 0 \end{pmatrix}, \quad (181)$$

$$A_\chi = L_{1,0} h_1^{-H^\dagger} \left[1 - \left(H^\dagger + \frac{1}{2} \right) \ell_1 \right], \quad (182)$$

$$B_\chi = \left(H^\dagger + \frac{1}{2} \right) h_1^{-H^\dagger}. \quad (183)$$

Then

$$J_0 = \begin{pmatrix} a_H' \\ (A_\chi a_H + B_\chi e_3)' \end{pmatrix}. \quad (184)$$

In particular,

$$\text{Avar}\{\sqrt{n}(\hat{H}_n - H^\dagger)\} = a'_H \Omega_0 a_H, \quad (185)$$

$$\text{Avar}\{\sqrt{n}(\hat{\chi}_{R,n} - \chi_R^\dagger)\} = A_\chi^2 a'_H \Omega_0 a_H + B_\chi^2 e'_3 \Omega_0 e_3 + 2A_\chi B_\chi a'_H \Omega_0 e_3. \quad (186)$$

Proof. Because $S_{1,0}, S_{2,0} > 0$, the map $m \mapsto g_H(m)$ is continuously differentiable in a neighborhood of m_0 . Consistency of $\hat{m}_n$ implied by Assumption 5.2 and the continuous mapping theorem give $\hat{H}_n \rightarrow_p H^\dagger$. The map

$$(H, L) \mapsto (H + \tfrac{1}{2})Lh_1^{-H}$$

is continuously differentiable on any compact subset of $(0, 1/2) \times \mathbb{R}$, hence $\hat{\chi}_{R,n} \rightarrow_p \chi_R^\dagger$.

Differentiation of (175) gives the vector a_H in (181). Next,

$$\frac{\partial}{\partial H} \left\{ (H + \tfrac{1}{2})Lh_1^{-H} \right\} = Lh_1^{-H} \left[1 - (H + \tfrac{1}{2}) \log h_1 \right],$$

whereas

$$\frac{\partial}{\partial L} \left\{ (H + \tfrac{1}{2})Lh_1^{-H} \right\} = (H + \tfrac{1}{2})h_1^{-H}.$$

The chain rule therefore gives the second row of (184). Applying the multivariate delta method to (174) yields (180). The marginal variance formulas follow by direct multiplication. $\square$

A useful feature of Theorem 5.3 is that $\hat{\chi}_{R,n}$ can be computed *directly* from $(\hat{H}_n, \hat{L}_{1,n})$. One need not first estimate $\hat{\lambda}_R$ and $\hat{\rho}_{SR}$ separately. This avoids unnecessary propagation of estimation noise when the object of interest is the leading rough ES correction.

Feasible covariance estimation. Let $\hat{\psi}_t$ be an estimated influence-function sequence and define

$$\hat{\Omega}_n = \hat{\Gamma}_0 + \sum_{k=1}^{b_n} w(k/b_n) (\hat{\Gamma}_k + \hat{\Gamma}'_k), \quad (187)$$

where w is a standard HAC kernel and $b_n \rightarrow \infty$, $b_n/n \rightarrow 0$. If $\hat{\Omega}_n \rightarrow_p \Omega_0$, then

$$\hat{V}_{H\chi} = \hat{J}_n \hat{\Omega}_n \hat{J}'_n \quad (188)$$

is consistent for $V_{H\chi}$ and the standard errors are $\sqrt{\hat{V}_{jj}/n}$.

5.3.1 Structural interpretation under shrinking pilot horizons

The fixed-horizon theorem estimates local effective parameters. To connect them rigorously to primitive rough parameters, we now state an abstract triangular-array implication. It is not, by itself, a high-frequency sampling theorem: a concrete infill application would also specify an observation mesh Δ_n , integer horizon lengths and the effective number of overlapping or non-overlapping blocks. Consider pilot horizons

$$h_{1,n} = b_n \downarrow 0, \quad h_{2,n} = r b_n, \quad r > 1. \quad (189)$$

Assume the local population moments admit the uniform expansions

$$\mathcal{S}_R(h) = A_0 h^{2H} \{1 + c_S h^\beta + o(h^\beta)\}, \quad A_0 = \frac{\lambda_R^2}{2H}, \quad (190)$$

$$\mathcal{L}_R(h) = B_0 h^H \{1 + c_L h^\beta + o(h^\beta)\}, \quad B_0 = \frac{\chi_R}{H + \frac{1}{2}}, \quad (191)$$

for some $\beta > 0$. These expansions allow for tempering and other smooth departures from an exact power-law kernel.

Assumption 5.4 (Relative triangular-array CLT). *With*

$$D_n = \text{diag}\{\mathcal{S}_R(h_{1,n}), \mathcal{S}_R(h_{2,n}), \mathcal{L}_R(h_{1,n})\},$$

and $\chi_R \neq 0$,

$$\sqrt{n} D_n^{-1} (\hat{m}_n - m_n) \implies N(0, \Omega_{\text{rel}}), \quad (192)$$

where m_n is the population moment vector at $(h_{1,n}, h_{2,n})$ and Ω_{rel} is finite.

Corollary 5.5 (Structural consistency and the logarithmic rate for $\hat{\chi}_R$). *Suppose (190)–(191) and Assumption 5.4 hold, and assume*

$$\frac{|\log b_n|}{\sqrt{n}} \longrightarrow 0. \quad (193)$$

Then

$$\hat{H}_n \xrightarrow{p} H, \quad \hat{\chi}_{R,n} \xrightarrow{p} \chi_R. \quad (194)$$

If, in addition,

$$\sqrt{n} b_n^\beta \longrightarrow 0, \quad (195)$$

then

$$\begin{pmatrix} \sqrt{n}(\hat{H}_n - H) \\ \frac{\sqrt{n}}{|\log b_n|}(\hat{\chi}_{R,n} - \chi_R) \end{pmatrix} \implies \begin{pmatrix} Z_H \\ \chi_R Z_H \end{pmatrix}, \quad (196)$$

where, for $Z = (Z_1, Z_2, Z_3)' \sim N(0, \Omega_{\text{rel}})$,

$$Z_H = \frac{Z_2 - Z_1}{2 \log r}. \quad (197)$$

Hence the structural estimator of χ_R is consistent but has the slower effective rate $\sqrt{n}/|\log b_n|$ when identification is forced to arbitrarily small horizons.

Proof. Using (190),

$$\frac{1}{2} \frac{\log\{\mathcal{S}_R(r b_n)/\mathcal{S}_R(b_n)\}}{\log r} = H + O(b_n^\beta),$$

so relative consistency of the two variance moments yields $\hat{H}_n - H = O_p(n^{-1/2}) + O(b_n^\beta)$. Similarly,

$$\hat{\chi}_{R,n} = (\hat{H}_n + \tfrac{1}{2}) \hat{L}_{1,n} b_n^{-\hat{H}_n}.$$

Define

$$u_{L,n} := \frac{\hat{L}_{1,n} - \mathcal{L}_R(b_n)}{\mathcal{L}_R(b_n)}, \quad c_n := \frac{1}{H + \frac{1}{2}} - \log b_n.$$

Since $b_n^{\hat{H}_n} = b_n^{-H} \exp\{-(\hat{H}_n - H) \log b_n\}$, a triangular Taylor expansion gives

$$\frac{\hat{\chi}_{R,n} - \chi_R}{\chi_R} = u_{L,n} + c_n(\hat{H}_n - H) + O_p(b_n^\beta) + O_p(a_n^2), \quad (198)$$

where

$$a_n := |c_n(\hat{H}_n - H)| + |u_{L,n}| + b_n^\beta. \quad (199)$$

Under Assumption 5.4, $u_{L,n} = O_p(n^{-1/2})$ and $\hat{H}_n - H = O_p(n^{-1/2}) + O(b_n^\beta)$. Condition (193), together with $b_n^\beta |\log b_n| \rightarrow 0$, makes the amplified estimation error and the quadratic remainder a_n^2 vanish in probability, proving consistency. Under (195), the centered limit of $\hat{H}_n$ is Z_H in (197). Dividing (198) by $|\log b_n|/\sqrt{n}$ makes the direct error in $\hat{L}_{1,n}$ and the quadratic remainder asymptotically negligible, while

$$\frac{(H + \frac{1}{2})^{-1} - \log b_n}{|\log b_n|} \rightarrow 1.$$

This yields (196). $\square$

Corollary 5.5 has an important practical implication. Pushing the pilot horizons ever closer to zero removes local-model bias but makes the structural amplitude–leverage product increasingly sensitive to estimation error in H . For empirical risk management, this argues for reporting both the structural $\hat{\chi}_R$ and the directly estimated horizon-specific leverage statistic $\hat{\mathcal{L}}_R(h)$.

5.4 Conditional delta-method inference for the rough ES shape correction

At a fixed forecast origin, the forward-variance scale $a_{t,h}$ is an $\mathcal{F}_t$ -measurable state variable. It need not converge to a deterministic constant when the historical sample length increases. We therefore formulate inference for the rough *shape* correction conditionally on the current state. Filtering uncertainty in $a_{t,h}$ is propagated separately by bootstrap or state-space simulation in the operational risk engine.

Fix h and $\alpha \in \mathcal{A} \subseteq (1/2, 1)$ and define

$$c_\alpha = \frac{\phi(z_\alpha)}{1 - \alpha}, \quad k_\alpha(H; h) := -\text{sgn}(\mathcal{E}_t) \frac{z_\alpha h^H}{2(H + \frac{1}{2})(H + \frac{3}{2})}. \quad (200)$$

The first-order normalized ES multiplier and conditional ES approximation are

$$\mathcal{M}_\alpha(H, \chi_R; h) := 1 + \chi_R k_\alpha(H; h), \quad (201)$$

$$\mathcal{E}_\alpha(H, \chi_R, a_{t,h}; h) := a_{t,h} c_\alpha \mathcal{M}_\alpha(H, \chi_R; h). \quad (202)$$

Set

$$q_H(h) := \log h - \frac{1}{H + \frac{1}{2}} - \frac{1}{H + \frac{3}{2}}. \quad (203)$$

Then

$$\nabla_{(H,\chi)} \mathcal{M}_\alpha = \begin{pmatrix} \chi_R k_\alpha(H; h) q_H(h) \\ k_\alpha(H; h) \end{pmatrix}, \quad (204)$$

and, conditional on $a_{t,h}$,

$$g_{ES,t} := \nabla_{(H,\chi)} \mathcal{E}_\alpha = a_{t,h} c_\alpha \begin{pmatrix} \chi_R k_\alpha(H; h) q_H(h) \\ k_\alpha(H; h) \end{pmatrix}. \quad (205)$$

Assumption 5.6 (Stable fixed-horizon pilot limit). *The fixed-horizon limit in Theorem 5.3 holds stably with respect to the forecast-origin sigma-field $\mathcal{F}_t$:*

$$\sqrt{n} \begin{pmatrix} \hat{H}_n - H^\dagger \\ \hat{\chi}_{R,n} - \chi_R^\dagger \end{pmatrix} \Rightarrow^{st} V_{H\chi}^{1/2} Z, \quad Z \sim N(0, I_2), \quad (206)$$

When the pilot moments are explicitly conditioned on the current state, $V_{H\chi}$ may be $\mathcal{F}_t$ -measurable; in that case Z is taken independent of the limiting $\mathcal{F}_t$ -field.

Theorem 5.7 (Conditional delta method for the rough ES shape). *Under Assumption 5.6, define*

$$\widehat{M}_{ES,\alpha,t}(h) := \mathcal{M}_\alpha(\hat{H}_n, \hat{\chi}_{R,n}; h), \quad (207)$$

$$\widehat{ES}_{\alpha,t}^R(h) := a_{t,h} c_\alpha \widehat{M}_{ES,\alpha,t}(h), \quad (208)$$

where $a_{t,h}$ is conditioned upon at the forecast origin. Then

$$\widehat{M}_{ES,\alpha,t}(h) \xrightarrow{p} \mathcal{M}_\alpha(H^\dagger, \chi_R^\dagger; h), \quad (209)$$

and

$$\sqrt{n} \left(\widehat{ES}_{\alpha,t}^R(h) - ES_{\alpha,t}^{R,\dagger}(h) \right) \Rightarrow^{st} MN(0, \sigma_{ES,t}^2), \quad (210)$$

where

$$ES_{\alpha,t}^{R,\dagger}(h) := a_{t,h} c_\alpha \mathcal{M}_\alpha(H^\dagger, \chi_R^\dagger; h) \quad (211)$$

and

$$\boxed{\sigma_{ES,t}^2 = g_{ES,t}^{\dagger'} V_{H\chi} g_{ES,t}^\dagger} \quad (212)$$

where $g_{ES,t}^\dagger$ denotes (205) evaluated at $(H^\dagger, \chi_R^\dagger)$. Conditionally on $\mathcal{F}_t$, the limit is Gaussian with variance $\sigma_{ES,t}^2$; unconditionally it is mixed normal whenever that variance is random.

Proof. The map $(H, \chi) \mapsto \mathcal{M}_\alpha(H, \chi; h)$ is continuously differentiable on compact subsets of $(0, 1/2) \times \mathbb{R}$, with gradient (204). The stable delta method applied to (206), followed by multiplication by the $\mathcal{F}_t$ -measurable factor $a_{t,h} c_\alpha$, proves (210)–(212). $\square$

A feasible conditional standard error is

$$\widehat{\text{se}}\{\widehat{ES}_{\alpha,t}^R(h)\} = \sqrt{\frac{\widehat{g}_{ES,t}^{\dagger'} \widehat{V}_{H\chi} \widehat{g}_{ES,t}^\dagger}{n}}. \quad (213)$$

This is a *shape-parameter* standard error conditional on the filtered scale. Total forecast

uncertainty also includes state-filtering error, model misspecification and Monte Carlo error. These components should be reported separately or propagated by a block or filtering bootstrap.

Remark 5.8 (Approximation error versus parameter-estimation error). The theorem concerns the first-order ES approximation. To use the same centered limit for the exact finite-horizon ES in a joint $n \rightarrow \infty$, $h = h_n \downarrow 0$ experiment, the analytical remainder from Section 4 must be negligible at the statistical scale. At fixed economically relevant horizons, delta-method shape uncertainty and full Monte Carlo revaluation should therefore be reported separately.

Remark 5.9 (Why inference on χ_R is preferable to separate inference on λ_R and ρ_{SR}). The first-order ES correction depends on λ_R and ρ_{SR} only through $\chi_R = \lambda_R \rho_{SR}$. Direct inference on χ_R therefore targets the parameter combination relevant for the risk measure and avoids unnecessary propagation of noise from a separate decomposition. Separate estimates remain useful for structural simulation and economic interpretation.

5.5 Stage 1: multiscale variogram and covariance estimation

The first-stage parameter vector is

$$\theta_V = (H, \sigma_R, \kappa_R, \sigma_D, \kappa_D, \rho_{DR}, \sigma_e^2)'. \quad (214)$$

We estimate θ_V from a collection of volatility moments spanning short, intermediate and long lags. Let

$$\hat{\gamma}_y(k) = \frac{1}{T-k} \sum_{t=1}^{T-k} (y_{t+k} - \bar{y})(y_t - \bar{y}) \quad (215)$$

and

$$\hat{V}_y(k) = \frac{1}{T-k} \sum_{t=1}^{T-k} (y_{t+k} - y_t)^2. \quad (216)$$

Because y_t contains $\log \int v_s ds$, its model moments are nonlinear functionals of the lognormal Volterra process. They are therefore evaluated by simulation, numerical quadrature, or a controlled approximation to the integrated-variance measurement equation. Closed-form kernel integrals are used for latent log-variance moments and initialization, but are not treated as exact moments of $\log IV_t$.

For a lag set $\mathcal{K} = \{k_1, \dots, k_J\}$, define the moment vector

$$g_T^V(\theta_V) = \begin{pmatrix} \hat{V}_y(k_1) - V_y(k_1; \theta_V) \\ \vdots \\ \hat{V}_y(k_J) - V_y(k_J; \theta_V) \\ \hat{\gamma}_y(k_1) - \gamma_y(k_1; \theta_V) \\ \vdots \\ \hat{\gamma}_y(k_J) - \gamma_y(k_J; \theta_V) \end{pmatrix}. \quad (217)$$

The preliminary estimator minimizes

$$\hat{\theta}_V^{(1)} \in \arg \min_{\theta_V \in \Theta_V} [g_T^V(\theta_V)]' W_V [g_T^V(\theta_V)], \quad (218)$$

where W_V is initially diagonal and subsequently replaced by a HAC estimate of the moment covariance matrix.

The lag grid is chosen to span the anticipated volatility crossover. Very short lags are primarily informative about the slope $2H$ and rough amplitude, intermediate lags are informative about curvature generated by the mixture, and long lags are informative about mean reversion. To diagnose weak separation of the two components, we report the condition number of the sample Jacobian

$$\widehat{G}_V = \frac{\partial g_T^V(\theta_V)}{\partial \theta_V'} \Big|_{\widehat{\theta}_V}. \quad (219)$$

A large condition number is interpreted as a local weak-identification diagnostic indicating that the data do not support a reliable rough–diffusive decomposition over the observed frequency range.

5.6 Stage 2: leverage and return dynamics

Conditional return–volatility dependence identifies

$$\theta_L = (\rho_{SR}, \rho_{SD}, \mu)'. \quad (220)$$

We use signed cross moments between returns and future changes in log realized variance,

$$\widehat{c}_{r,v}(k) = \frac{1}{T-k} \sum_{t=1}^{T-k} (r_t - \bar{r})(y_{t+k} - y_t), \quad k \in \mathcal{K}_L, \quad (221)$$

together with return variance and selected higher moments. The model counterparts are obtained by simulation conditional on the Stage-1 volatility parameters. The second-stage criterion is

$$\widehat{\theta}_L \in \arg \min_{\theta_L \in \Theta_L} [g_T^L(\widehat{\theta}_V, \theta_L)]' W_L [g_T^L(\widehat{\theta}_V, \theta_L)]. \quad (222)$$

This separation reduces the tendency for leverage parameters to compensate mechanically for an incorrectly estimated volatility spectrum; it does not eliminate such compensation in the final joint stage.

5.7 Stage 3: joint simulated minimum distance

The final estimator refines all structural parameters jointly. Let

$$\theta = (H, \sigma_R, \kappa_R, \sigma_D, \kappa_D, \rho_{DR}, \rho_{SR}, \rho_{SD}, \bar{v}, \mu, \sigma_e^2)'. \quad (223)$$

We augment the volatility and leverage moments with direct-horizon return moments at $h \in \mathcal{H}_E$, including variance, downside semivariance, skewness, excess kurtosis and selected quantiles. Denote the complete vector by $g_T(\theta)$. Simulated minimum distance uses

$$\widehat{\theta} \in \arg \min_{\theta \in \Theta} g_T(\theta)' \widehat{W}_T g_T(\theta), \quad (224)$$

where $\widehat{W}_T$ is a regularized HAC weighting matrix. Common random numbers are used across nearby parameter values to reduce simulation noise in the criterion.

The direct inclusion of multi-horizon tail moments is intentionally limited. We do not estimate the model by forcing it to reproduce the same VaR and ES forecasts that will later be used for evaluation. The baseline criterion therefore includes only a small number of distributional moments; out-of-sample VaR and ES remain genuine forecast targets.

5.8 Testing roughness and the need for a second scale

The empirical analysis treats the mixed-rough specification as a hypothesis rather than as a maintained truth. We consider the nested restrictions

$$H_0^{R0} : \sigma_R = 0, \quad (225)$$

$$H_0^{D0} : \sigma_D = 0, \quad (226)$$

$$H_0^M : H = \frac{1}{2}, \quad (227)$$

and compare the full model with single-factor rough and two-factor Markov alternatives. Because some restrictions place parameters on a boundary or change the effective dimension of the model, standard likelihood-ratio asymptotics need not apply. We therefore combine information criteria with a parametric bootstrap of the distance-statistic difference and, most importantly, out-of-sample forecast comparison.

To address the possibility of spurious roughness, four robustness checks are mandatory:

1. repeat estimation with alternative realized-volatility measures and intraday sampling frequencies;
2. include an explicit measurement-error variance σ_e^2 ;
3. estimate the short-lag slope after omitting the first one or two lags;
4. simulate data from the fitted non-rough benchmark, pass them through the same realized-volatility measurement procedure, and measure the distribution of the pseudo-estimated $\widehat{H}$.

The final comparison asks whether roughness remains distinguishable after the measurement-error experiment.

5.9 Rolling estimation for risk forecasting

Risk forecasts are produced recursively or with a fixed rolling window. At each forecast origin t , parameters are estimated using information available through t , yielding $\widehat{\theta}_t$. For each horizon h and confidence level α , the model generates

$$\widehat{\text{VaR}}_{\alpha,t}(h), \quad \widehat{\text{ES}}_{\alpha,t}(h), \quad \widehat{\text{RHM}}_{\alpha,t}(h \mid h_0). \quad (228)$$

Parameters that govern very persistent features, such as H , may be updated less frequently than the current variance state. The baseline therefore compares two implementations: full rolling

re-estimation and a hybrid scheme in which structural parameters are updated monthly while the latent state and current forward variance are updated daily.

5.10 Separating physical and risk-neutral estimation

For cash portfolios, the physical estimator above is sufficient. For derivative portfolios, we additionally estimate or calibrate a risk-neutral map

$$\Theta_t^{\mathbb{Q}} = \mathcal{M}(\theta_t^{\mathbb{P}}, \Lambda_t), \quad (229)$$

where Λ_t denotes volatility-risk-premium parameters. The baseline keeps $\mathcal{M}$ parsimonious so that tail forecasts are not driven by overfitting the option surface. A richer state-dependent volatility-risk-premium specification is used as a robustness check, following the conceptual separation emphasized in Bonesini et al. (2025).

The estimation section therefore delivers not only a point estimate of H , but a complete multi-scale decomposition and an identification diagnostic. Those quantities are carried forward into the risk engine developed next.

6 Direct-Horizon Risk Engine, Benchmarks, and Backtesting

This section describes the numerical procedure used to transform the estimated physical- measure model into horizon-specific portfolio-loss distributions. The emphasis is on direct simulation at the target holding period rather than on ex post scaling of a shorter-horizon risk number.

6.1 Simulation of the mixed rough volatility factors

The diffusive factor is simulated exactly on the observation grid. For step size Δ ,

$$X_{t+\Delta}^D = e^{-\kappa_D \Delta} X_t^D + \sqrt{1 - e^{-2\kappa_D \Delta}} \varepsilon_{t+\Delta}^D, \quad (230)$$

with standard normal innovations appropriately correlated with the return and rough-factor innovations.

The rough factor is simulated using a hybrid convolution scheme in the spirit of Bennedsen et al. (2017). The kernel singularity is treated analytically over the nearest grid cells while the smooth tail is discretized by a fast convolution. This avoids the large bias that a naive Euler discretization produces when H is small. For long horizons or repeated rolling forecasts, the convolution is implemented with FFT acceleration.

At each Monte Carlo path, the variance is recovered from

$$v_{t_j} = \bar{v} \exp \left(Y_{t_j} - \frac{1}{2} \omega_Y^2 \right), \quad (231)$$

and integrated variance is approximated by a trapezoidal or higher-order quadrature. The price process is then advanced under $\mathbb{P}$ with correlated innovations consistent with (7).

6.2 Direct-horizon VaR and ES

For forecast origin t and target horizon h , simulate N conditional losses

$$L_{t,h}^{(1)}, \dots, L_{t,h}^{(N)}. \quad (232)$$

Let $L_{(1)} \leq \dots \leq L_{(N)}$ denote their order statistics and $k_\alpha = \lceil N\alpha \rceil$. The Monte Carlo estimators are

$$\widehat{\text{VaR}}_{\alpha,t}^{MC}(h) = L_{(k_\alpha)}, \quad (233)$$

$$\widehat{\text{ES}}_{\alpha,t}^{MC}(h) = \frac{1}{N - k_\alpha + 1} \sum_{j=k_\alpha}^N L_{(j)}. \quad (234)$$

For high confidence levels, tail Monte Carlo uncertainty is reported explicitly. The baseline uses independent batch means to estimate the simulation standard error of ES and increases N until simulation error is small relative to the empirical forecast error.

6.3 Full revaluation of derivatives

For an option portfolio, each physical path generates the horizon state $(S_{t+h}, v_{t+h}, \dots)$. The book is then repriced under the risk-neutral mapping (229). To avoid an infeasible nested simulation at every outer path, we use three increasingly accurate valuation layers:

1. a Delta–Gamma approximation;
2. a Delta–Gamma–Vega–Vanna–Volga approximation;
3. a full-revaluation benchmark using either a semi-analytic rough-volatility pricing routine, a precomputed pricing surface, or a surrogate model trained on risk-neutral prices.

Common random numbers and antithetic sampling are used wherever compatible with the pricing routine. The gap between the Greek approximation and full revaluation is itself treated as a model-risk quantity.

6.4 Benchmark risk models

The mixed-rough model is evaluated against a deliberately broad but interpretable benchmark set:

1. historical simulation;
2. Gaussian parametric VaR/ES;
3. Student- t parametric VaR/ES;
4. Cornish–Fisher VaR with its corresponding moment-based ES approximation;
5. EWMA;
6. GARCH- t ;

7. a one-factor diffusive stochastic-volatility model;
8. Heston;
9. a single-factor rough volatility model;
10. a two-factor non-rough Markov stochastic-volatility model.

The last benchmark is particularly important. If the mixed-rough model performs better than Heston but not better than a flexible two-factor Markov model, the improvement cannot be attributed uniquely to roughness.

6.5 Conventional and regulatory horizon benchmarks

For each model we distinguish direct-horizon risk from transported risk. Given reference horizon h_0 ,

$$\widehat{ES}_{\alpha,t}^{\sqrt{\cdot}}(h \mid h_0) = \widehat{ES}_{\alpha,t}(h_0)\sqrt{h/h_0}. \quad (235)$$

The corresponding empirical RHM is computed from (45).

For the FRTB comparison, we do not replace the regulatory construction by a simple one-day-to-ten-day square-root rule. Let $LH_1 = 10 < LH_2 < \dots < LH_J$ denote the prescribed liquidity horizons and let $ES_T(P, j)$ denote the ten-day ES of the portfolio after setting to zero the risk factors whose liquidity horizon is shorter than LH_j . The regulatory aggregation has the schematic form

$$ES^{LH}(P) = \sqrt{ES_T(P, 1)^2 + \sum_{j=2}^J ES_T(P, j)^2 \frac{LH_j - LH_{j-1}}{T}}, \quad T = 10, \quad (236)$$

with the exact implementation following the applicable Basel specification (Basel Committee on Banking Supervision, 2020). Our comparison is between this liquidity-horizon architecture and the direct model-implied loss distribution over the corresponding economic holding periods.

6.6 VaR backtesting

For each horizon, VaR violations are

$$I_{t+h}(\alpha, h) = \mathbf{1}\{L_{t,t+h} > \widehat{\text{VaR}}_{\alpha,t}(h)\}. \quad (237)$$

The evaluation reports unconditional exception frequency, independence and conditional coverage diagnostics, including the framework of Christoffersen (1998). Because multi-day forecasts overlap when produced daily, inference uses either non-overlapping forecast origins or heteroskedasticity-and-autocorrelation robust standard errors, with the non-overlapping design as the primary robustness benchmark.

6.7 Expected-Shortfall backtesting and joint scoring

Exception counts alone cannot assess the severity of tail losses. We therefore complement VaR tests with ES backtests, including the conditional calibration approach in the spirit of Du and

Escanciano (2017). Model ranking is based primarily on strictly consistent joint VaR–ES scores from the class characterized by Fissler and Ziegel (2016). Let

$$S_\alpha(q, e; L) \quad (238)$$

be a member of this class, where q is VaR and e is ES under the paper’s positive-loss convention. Average out-of-sample score differences are compared with HAC inference to account for serial dependence and overlapping horizons.

The evaluation hierarchy is therefore:

$$\text{calibration} \implies \text{joint VaR–ES score} \implies \text{economic capital impact}. \quad (239)$$

A model is not considered useful merely because it produces a larger ES. It must improve calibration or predictive scores, and the improvement must be large enough to matter relative to estimation and simulation uncertainty.

6.8 Risk-attribution outputs for practitioners

At each forecast origin the risk engine stores the following objects:

$$\{VaR_\alpha(h), ES_\alpha(h), RHM_\alpha(h \mid h_0), HTA_\alpha(h \mid h_0), RRA_\alpha(h), \beta_\alpha(h)\}_{h \in \mathcal{H}}. \quad (240)$$

For derivative books, these are complemented by position-level ES contributions and by the Greek/full-revaluation gap. The purpose of these quantities is to turn an abstract rough-volatility parameter into operational questions: at which horizon does the conventional transport rule become inaccurate, how large is the monetary error, and which positions or risk sensitivities generate it?

7 Monte Carlo Validation and Extended Design

The Monte Carlo analysis has a deliberately narrow first objective: to check that the local coefficients derived in Section 4 describe the finite-horizon loss simulation with the expected sign, rate and magnitude. A second experiment evaluates the closed-form pilot estimators and the delta-method standard errors from Section 5. These exercises are theorem-validation diagnostics rather than empirical evidence in favor of rough volatility. A broader model-comparison design is then specified for the empirical stage.

7.1 Validation of the rough-leverage ES expansion

To isolate the mechanism in Theorem 4.22, we suppress the diffusive factor and use a local rough log-volatility experiment with flat conditional forward variance. On the unit interval let

$$Y_u = \int_0^u (u-s)^{H-1/2} dB_s, \quad d\langle W, B \rangle_u = \rho du, \quad (241)$$

and, for a physical holding horizon h ,

$$v_{hu} = \exp \left\{ \lambda h^H Y_u - \frac{1}{2} \lambda^2 h^{2H} \text{Var}(Y_u) \right\}. \quad (242)$$

The compensation makes $\mathbb{E}[v_{hu}] = 1$. For a long linear position, the normalized martingale loss is

$$\tilde{L}_h = - \int_0^1 \sqrt{v_{hu}} dW_u. \quad (243)$$

This representation follows from Brownian scaling and has matched-Gaussian variance close to one. The experiment uses

$$H = 0.12, \quad \lambda = 0.09512, \quad \rho = -0.70, \quad \alpha = 0.975. \quad (244)$$

We simulate 120,000 paths on 256 time steps using cell-average quadrature for the singular kernel and seed 20260924. The same normalized Brownian paths are reused across horizons, so differences across h are driven by the factor h^H rather than by independent Monte Carlo noise. Batch means are used to monitor simulation uncertainty.

For these parameters, Theorem 4.9 gives

$$\mathfrak{s}_R = - \frac{3\lambda\rho}{(H + \frac{1}{2})(H + \frac{3}{2})} = 0.1988, \quad (245)$$

so that the predicted standardized skewness is $\mathfrak{s}_R h^H$. Table 1 compares the full nonlinear simulation with the first-order formulas. The last column is the percentage error of the analytical ES relative to the simulated ES.

Table 1: Monte Carlo validation of the rough-leverage VaR and ES expansion

h (days)	Skew MC	Skew asym.	VaR MC	VaR asym.	ES MC	ES asym.	ES error (%)
1	0.1004	0.1024	2.0084	2.0085	2.4296	2.4160	-0.56
2	0.1094	0.1113	2.0122	2.0127	2.4369	2.4228	-0.58
5	0.1225	0.1242	2.0195	2.0188	2.4475	2.4327	-0.61
10	0.1334	0.1350	2.0241	2.0239	2.4565	2.4409	-0.63
20	0.1452	0.1467	2.0278	2.0295	2.4662	2.4499	-0.66
40	0.1581	0.1595	2.0364	2.0355	2.4767	2.4596	-0.69

Notes: VaR and ES are reported in matched-standard-deviation units at the 97.5% confidence level. “Asym.” denotes the first-order Edgeworth approximation from Section 4. Monte Carlo uses 120,000 paths and 256 time steps.

The log–log slope of simulated skewness over the reported horizons is 0.1232, close to the structural value $H = 0.12$. The VaR approximation is already very accurate. For ES, the first-order expression understates the full nonlinear Monte Carlo value slightly, but the relative discrepancy remains below 0.7% from one to 40 trading days. The difference is consistent with higher-order rough terms and time-grid approximation. For scale, a EUR 10 million linear equity book with 1.2% daily volatility has a matched-Gaussian 97.5% ten-day ES of approximately EUR 0.887 million; the first-order rough-leverage multiplier raises this to about EUR 0.926 million, while the nonlinear Monte Carlo value is about EUR 0.932 million. These figures are illustrative rather than empirical. The experiment supports the coefficient and the h^H rate

used in the paper; it does not replace the Fourier remainder condition required for the uniform Edgeworth theorem.




Figure 1: Monte Carlo and first-order rough-leverage ES shape add-ons relative to a matched Gaussian benchmark.

7.2 Finite-sample recovery and ES-multiplier inference

We next validate the pilot estimator in a controlled experiment where the local moment restrictions hold exactly. Let

$$\chi_R = \lambda\rho = -0.066584, \quad h_1 = 1/252, \quad h_2 = 16/252. \quad (246)$$

For each replication we generate $n = 2,000$ Gaussian observations with the exact covariance matrix implied by the rough innovations at h_1 and h_2 and by their return–volatility cross moment. We then compute $\hat{H}$ and $\hat{\chi}_R$ from (177)–(178). There are 5,000 independent replications. The ten-day ES object is the shape multiplier

$$M_{ES}(10) = 1 + \chi_R k_{0.975}(H; 10/252), \quad (247)$$

whose true value is 1.044108.

Table 2: Parameter recovery and delta-method inference

Quantity	True	MC mean	MC SD	Delta SE	95% coverage
H	0.12000	0.11989	0.00794	0.00803	95.56
χ_R	-0.06658	-0.06658	0.00358	0.00358	95.22
$M_{ES}(10)$	1.04411	1.04411	0.00183	0.00183	94.58

Notes: $n = 2,000$ observations per replication and 5,000 replications. Delta SE is the asymptotic standard error implied by Theorem 5.3 and Theorem 5.7. Coverage uses two-sided normal 95% intervals.

The simulation delivers three checks. First, bias in both structural quantities is negligible at $n = 2,000$. Second, the analytical standard deviations are almost identical to the Monte Carlo

standard deviations. Third, the nominal 95% intervals attain coverage between 94.6% and 95.6%. In particular, the delta method transmits parameter uncertainty to the economically interpretable ES multiplier without a visible finite-sample distortion in this benchmark experiment.

7.3 Extended simulation design for model discrimination

The validation exercises above isolate mechanisms for which the theory is explicit. The full numerical study uses the stationary mixed model from Section 3 and asks a broader question: can the additional rough factor be distinguished from flexible non-rough volatility in samples of realistic length? The principal grid varies

$$H \in \{0.05, 0.10, 0.20, 0.30, 0.40\}, \quad \rho_{SR} \in \{-0.9, -0.7, -0.4, 0\}, \quad (248)$$

with low, medium and high rough-factor variance shares. Mean-reversion parameters are chosen so that the volatility crossover lies below five days, around 10–20 days, or beyond 40 days. Each data-generating process is fitted by a one-factor Markov model, a two-factor Markov model, a one-factor rough model and the mixed-rough specification. Reverse experiments generate data from the two-factor Markov model to measure pseudo-roughness.

A second block adds noise to realized-volatility proxies and compares estimation with and without the measurement equation. A third block revalues stylized option books with Delta, Gamma, Vega and cross-Greek exposures. The central outputs are parameter bias and RMSE, coverage, joint VaR–ES scores, horizon-scaling distortion, and the gap between Greek approximations and full revaluation. These experiments are kept separate from the reported validation tables because their role is model discrimination rather than verification of the local analytical coefficients.

7.4 Reproducibility

All reported Monte Carlo tables are generated from stored numerical output. The validation experiment uses a fixed seed, records the discretization grid and reports simulation standard errors. The empirical replication package will use the same principle for the full mixed model: parameters, random-number seeds, discretization choices and pricing routines are archived so that numerical error can be distinguished from economic model error.

8 Empirical and Practitioner Validation Protocol

This section specifies a prospective empirical validation protocol for a narrow question: does the mixed rough mechanism improve the prediction of the horizon profile of tail losses, and if so, for which portfolios and market states? The analysis therefore prioritizes a liquid equity index and its listed options, where high-frequency volatility measures and option surfaces can be combined consistently. A fixed-income extension is left to a robustness appendix so that the main paper remains focused.

8.1 Data architecture

A full empirical implementation would combine four data layers:

1. daily index prices and total returns;
2. intraday prices used to construct realized variance and alternative noise-robust volatility measures;
3. option quotes used to construct the implied-volatility surface and to revalue the representative option books;
4. short-rate or discount-curve inputs needed for option pricing.

The final sample dates should be chosen according to data availability, and the evaluation window should contain both calm and stressed periods and must be long enough to support genuine rolling out-of-sample estimation.

Cleaning rules should be fixed before forecast evaluation. Option filters remove stale quotes, violations of elementary arbitrage bounds, extremely small time value and contracts with insufficient liquidity. Intraday data are sampled at several frequencies to assess the sensitivity of roughness estimates to microstructure noise.

8.2 Representative portfolios

The first portfolio is a linear equity book whose value moves one-for-one with the index. It provides the cleanest comparison with the local theory of Section 4. The second portfolio is a diversified cash-equity book with time-varying weights but no options. The third is a stylized equity-options book containing calls and puts across moneyness and maturity buckets. The option book is chosen so that Delta, Gamma, Vega, Vanna and Volga contributions are all economically non-negligible.

For the options book, the position set is held fixed over each forecast interval and is rebalanced only at predetermined dates. This avoids confounding forecast performance with endogenous trading. The paper reports both unhedged and Delta-hedged versions because rough-volatility risk should be more visible once first-order spot exposure has been reduced.

8.3 Forecast origins and horizons

At each forecast origin t , all competing models are estimated using exactly the same information set. The horizon grid is

$$\mathcal{H}_E = \{1, 5, 10, 20, 40, 60, 120\} \text{ trading days,} \quad (249)$$

with primary emphasis on 10, 20, 40, 60 and 120 days for liquidity-horizon interpretation. The confidence levels are

$$\alpha \in \{0.95, 0.975, 0.99\}. \quad (250)$$

The 97.5% level is the principal ES specification, while 99% is included because of its widespread use in internal risk management.

Because daily forecasts at multi-day horizons overlap, the main statistical comparison uses non-overlapping forecast origins for formal tests and the full daily sequence for descriptive dashboard plots. HAC inference is reported as a secondary approach on the overlapping sample.

8.4 Competing models

All benchmarks in Section 6 are re-estimated in the same rolling scheme. The central comparison is not mixed rough versus Gaussian; it is mixed rough versus the strongest plausible alternatives:

$$\text{GARCH-}t, \quad \text{Heston}, \quad \text{two-factor Markov SV}, \quad \text{single-factor rough SV}, \quad \text{mixed rough SV}. \quad (251)$$

Simpler Normal, Student and Cornish–Fisher methods remain in the tables because they are common practitioner benchmarks and make the incremental computational cost transparent.

8.5 Forecast evaluation

For each model m , horizon h and confidence level α , we obtain

$$(\hat{q}_{m,t}^{\alpha,h}, \hat{e}_{m,t}^{\alpha,h}) = (\widehat{\text{VaR}}_{m,\alpha,t}(h), \widehat{\text{ES}}_{m,\alpha,t}(h)). \quad (252)$$

Evaluation has four layers.

First, VaR calibration is assessed using exception rates and conditional coverage. Second, ES calibration is assessed using tail-residual diagnostics and ES backtests. Third, models are compared using strictly consistent joint VaR–ES scores. Fourth, economic significance is measured by the monetary forecast error and by the stability of the risk signal across horizons.

For model pairs m and b , define the loss differential

$$d_t^{m,b}(\alpha, h) = S_\alpha(\hat{q}_{m,t}, \hat{e}_{m,t}; L_{t,t+h}) - S_\alpha(\hat{q}_{b,t}, \hat{e}_{b,t}; L_{t,t+h}). \quad (253)$$

Negative average $d_t^{m,b}$ favors model m . Statistical significance is assessed with serial-correlation-robust inference and, where appropriate, block bootstrap confidence intervals.

8.6 State dependence and stress periods

A key empirical prediction is that horizon distortion should be state dependent. We therefore classify forecast origins using observable volatility and market-stress indicators. The baseline split uses terciles of current realized volatility; robustness checks use option- implied volatility and drawdown states. For each state s , we report

$$\mathbb{E}[\text{RHM}_{\alpha,t}(h \mid h_0) \mid s_t = s] \quad (254)$$

and the corresponding out-of-sample score differential.

The paper does not define stressed dates ex post from model failures. Stress windows are pre-specified using market-wide events or objective volatility thresholds. This avoids selecting episodes in which the rough model happens to look favorable.

8.7 Practitioner dashboard

The intended empirical output is organized around a compact risk dashboard. At each date it would report

$$(ES_{97.5\%,10d}, RHM_{97.5\%,20d}, RHM_{97.5\%,40d}, h^*, RRA_{97.5\%,10d}, HTA_{97.5\%,40d}), \quad (255)$$

plus the leading ES contributor for the option portfolio. A typical interpretation is not “rough volatility adds risk” but rather “the model implies that direct 40-day ES is materially above the amount obtained by transporting 10-day ES, and most of the difference comes from volatility-sensitive option positions.”

8.8 Economic materiality criterion

To guard against statistically significant but economically trivial improvements, we define a materiality threshold relative to portfolio value or risk budget. Let

$$M_{\alpha,t}(h) = \frac{|ES_{\alpha,t}^{MR}(h) - ES_{\alpha,t}^B(h)|}{\Pi_t}, \quad (256)$$

where B is the strongest non-rough benchmark. The mixed rough model is considered practically relevant only when better predictive performance is accompanied by a persistent and economically meaningful $M_{\alpha,t}(h)$.

8.9 What would constitute evidence in favor of the model?

The empirical burden of proof is intentionally high. Evidence in favor of the mixed rough specification requires three findings jointly: (i) roughness and a second volatility scale remain identified after measurement-error corrections, (ii) direct-horizon VaR–ES forecasts improve out-of-sample relative to strong non-rough alternatives, and (iii) the improvement is concentrated at horizons or portfolio exposures for which the theory predicts a material rough correction. Failure of any one of these conditions weakens the economic case for the additional model complexity.

9 Discussion and Conclusion

This paper studies a question that is simpler to state than to answer: how quickly does tail risk accumulate when the holding period of a portfolio increases? The conventional square-root-of-time rule provides an exact answer only under restrictive distributional and dependence assumptions. Rough-volatility models, on the other hand, describe a volatility process with unusually irregular short-scale dynamics, but roughness of volatility does not imply that a continuous semimartingale return itself scales as h^H . The appropriate comparison is therefore not h^H versus $\sqrt{h}$ for returns. It is the size and economic importance of the higher-order corrections to the local $\sqrt{h}$ tail-risk benchmark.

The model developed in this paper combines a diffusive mean-reverting volatility factor with a tempered rough Volterra factor under the physical measure. The mixture produces a

volatility crossover: rough innovations dominate sufficiently short volatility changes, whereas mean reversion and the diffusive component become increasingly important at longer horizons. This structure is then transmitted to portfolio losses through random integrated variance, return–volatility leverage and nonlinear revaluation.

The theory yields a particularly transparent local implication. For a locally linear portfolio, VaR and Expected Shortfall retain their leading $\sqrt{h}$ scaling. When the rough volatility shock is negatively correlated with the asset return, however, the standardized loss distribution acquires a positive skewness correction of order h^H . Expected Shortfall therefore admits an expansion of the form

$$ES_{\alpha,t}(h) = C_{\alpha,t}\sqrt{h} \left[1 + A_{\alpha,t}h^H + o(h^H) \right] \quad (257)$$

in the rough-leverage case. If the rough leverage channel is absent, random integrated variance produces a positive excess-kurtosis correction of relative order h^{2H} . These results provide a structural interpretation of horizon-scaling error without postulating a new universal power law for returns.

The paper translates this theory into two practitioner quantities. The Rough Horizon Multiplier compares direct-horizon ES with square-root transport from a reference horizon, while the Horizon-Transport Add-on expresses the same discrepancy in currency units. These measures are model-risk diagnostics rather than regulatory substitutes. In particular, the paper keeps the Basel liquidity-horizon construction conceptually separate from direct simulation and asks how far the two may diverge under an empirically estimated volatility process.

For nonlinear books, the relevant mechanism is richer. Delta, Gamma, Vega, Vanna and Volga exposures create several channels through which a change in the volatility state affects the loss distribution. This is why the paper compares Greek-based approximations with full revaluation rather than treating the option book as a linear exposure with an adjusted volatility. The resulting decomposition is intended to tell a risk manager not only whether ES changes under mixed rough volatility, but which positions and sensitivities are responsible for the change.

The proposed empirical validation protocol deliberately subjects the model to a demanding comparison. Roughness would be estimated under the physical measure, with explicit controls for measurement error, and the full specification would be compared with GARCH- t , Heston, a flexible two-factor Markov stochastic-volatility model and a one-factor rough model. VaR and ES are evaluated jointly out of sample using calibration diagnostics, ES backtests and strictly consistent scoring rules. The model is economically relevant only if any statistical gain is accompanied by a material improvement in the measurement of multi-horizon portfolio risk.

The numerical evidence reported in this version is deliberately limited to controlled Monte Carlo validation of the analytical coefficients and the pilot-estimator inference. No external market-data horse race is reported here; Section 8 is a pre-specified implementation and validation protocol intended to make subsequent empirical testing transparent and falsifiable.

Several extensions are natural. First, the volatility mixture can be embedded in a richer multivariate setting in which equity, rates and credit spreads possess distinct rough and mean-reverting components. Second, stochastic liquidity horizons can be added so that both the time required to unwind a position and the distribution of losses during that time are random. Third, the direct-horizon ES framework can be combined with optimal hedging or capital

allocation, making the roughness-related add-on an endogenous input to portfolio decisions. Finally, a separate treatment of jumps would allow the model to distinguish the horizon effects of discontinuous price risk from those generated by continuous but rough volatility.

The broader message is that tail risk has a term structure. A credible risk model should therefore explain not only the level of one-day or ten-day Expected Shortfall, but also the mechanism by which the loss distribution changes as the horizon used to hold, hedge or unwind a portfolio increases. Mixed rough volatility offers one structural explanation. Whether that explanation earns its additional complexity is ultimately an empirical question—and that is the question the proposed forecasting and backtesting framework is designed to answer.

Data and code availability

No external empirical data are used in the numerical results reported in this manuscript. The reported numerical evidence is generated by Monte Carlo simulation. Reproducible code, random-number seeds and the numerical outputs underlying the validation tables are supplied with the submission package.

References

- Abi Jaber, Eduardo, Martin Larsson, and Sergio Pulido (2019). “Affine Volterra Processes”. In: *The Annals of Applied Probability* 29.5, pp. 3155–3200. DOI: 10.1214/19-AAP1477.
- Acerbi, Carlo and Dirk Tasche (2002). “On the Coherence of Expected Shortfall”. In: *Journal of Banking & Finance* 26.7, pp. 1487–1503. DOI: 10.1016/S0378-4266(02)00283-2.
- Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber, and David Heath (1999). “Coherent Measures of Risk”. In: *Mathematical Finance* 9.3, pp. 203–228. DOI: 10.1111/1467-9965.00068.
- Basel Committee on Banking Supervision (2020). *Basel Framework: Market Risk, MAR33 — Internal Models Approach: Capital Requirements Calculation*. Bank for International Settlements. Basel Framework, MAR33; effective 1 January 2023; accessed September 2026. URL: <https://www.bis.org/committees/bcbs/basel-framework/standard/mar/33/inforce/2023-01-01/published/2020-06-05>.
- Bayer, Christian, Peter Friz, and Jim Gatheral (2016). “Pricing under Rough Volatility”. In: *Quantitative Finance* 16.6, pp. 887–904. DOI: 10.1080/14697688.2015.1099717.
- Bayer, Christian, Peter K. Friz, Archil Gulisashvili, Blanka Horvath, and Benjamin Stemper (2019). “Short-Time Near-the-Money Skew in Rough Fractional Volatility Models”. In: *Quantitative Finance* 19.5, pp. 779–798. DOI: 10.1080/14697688.2018.1529420.
- Bennedsen, Mikkel, Asger Lunde, and Mikko S. Pakkanen (2017). “Hybrid Scheme for Brownian Semistationary Processes”. In: *Finance and Stochastics* 21.4, pp. 931–965. DOI: 10.1007/s00780-017-0335-5.
- (2022). “Decoupling the Short- and Long-Term Behavior of Stochastic Volatility”. In: *Journal of Financial Econometrics* 20.5, pp. 961–1006. DOI: 10.1093/jjfinec/nbaa049.

- Bolko, Anine E., Kim Christensen, Mikko S. Pakkanen, and Bezirgen Veliyev (2023). “A GMM Approach to Estimate the Roughness of Stochastic Volatility”. In: *Journal of Econometrics* 235.2, pp. 745–778. DOI: 10.1016/j.jeconom.2022.06.009.
- Bollerslev, Tim (1986). “Generalized Autoregressive Conditional Heteroskedasticity”. In: *Journal of Econometrics* 31.3, pp. 307–327. DOI: 10.1016/0304-4076(86)90063-1.
- Bonesini, Ofelia, Antoine Jacquier, and Aitor Muguruza (2025). “Risk Premium and Rough Volatility”. In: *Frontiers of Mathematical Finance* 7, pp. 78–94. DOI: 10.3934/fmf.2025013.
- Bourgey, Florian, Stefano De Marco, and Emmanuel Gobet (2023). “Weak Approximations and VIX Option Price Expansions in Forward Variance Curve Models”. In: *Quantitative Finance* 23.9, pp. 1259–1283. DOI: 10.1080/14697688.2023.2227230.
- Cai, Zongwu and Xian Wang (2008). “Nonparametric Estimation of Conditional VaR and Expected Shortfall”. In: *Journal of Econometrics* 147.1, pp. 120–130. DOI: 10.1016/j.jeconom.2008.09.005.
- Cheridito, Patrick (2001). “Mixed Fractional Brownian Motion”. In: *Bernoulli* 7.6, pp. 913–934. DOI: 10.2307/3318626.
- Christoffersen, Peter F. (1998). “Evaluating Interval Forecasts”. In: *International Economic Review* 39.4, pp. 841–862. DOI: 10.2307/2527341.
- Colldeforns-Papiol, Gemma and Luis Ortiz-Gracia (2018). “Computation of Market Risk Measures with Stochastic Liquidity Horizon”. In: *Journal of Computational and Applied Mathematics* 342, pp. 431–450. DOI: 10.1016/j.cam.2018.03.038.
- Comte, Fabienne and Eric Renault (1998). “Long Memory in Continuous-Time Stochastic Volatility Models”. In: *Mathematical Finance* 8.4, pp. 291–323. DOI: 10.1111/1467-9965.00057.
- Cont, Rama and Purba Das (2024). “Rough Volatility: Fact or Artefact?” In: *Sankhya B* 86.1, pp. 191–223. DOI: 10.1007/s13571-024-00322-2.
- Daniélsson, Jón and Jean-Pierre Zigrand (2006). “On Time-Scaling of Risk and the Square-Root-of-Time Rule”. In: *Journal of Banking & Finance* 30.10, pp. 2701–2713. DOI: 10.1016/j.jbankfin.2005.10.002.
- Djeutcha, Eric and Jules Sadefo Kamdem (2021). “Local and Implied Volatilities with the Mixed-Modified-Fractional-Dupire Model”. In: *Chaos, Solitons & Fractals* 152, p. 111328. DOI: 10.1016/j.chaos.2021.111328.
- Drost, Feike C. and Theo E. Nijman (1993). “Temporal Aggregation of GARCH Processes”. In: *Econometrica* 61.4, pp. 909–927. DOI: 10.2307/2951767.
- Du, Zaichao and Juan Carlos Escanciano (2017). “Backtesting Expected Shortfall: Accounting for Tail Risk”. In: *Management Science* 63.4, pp. 940–958. DOI: 10.1287/mnsc.2015.2342.
- Dupret, Jean-Loup and Donatien Hainaut (2021). “Portfolio Insurance under Rough Volatility and Volterra Processes”. In: *International Journal of Theoretical and Applied Finance* 24.6–7, p. 2150036. DOI: 10.1142/S0219024921500369.
- El Euch, Omar, Masaaki Fukasawa, Jim Gatheral, and Mathieu Rosenbaum (2019). “Short-Term At-the-Money Asymptotics under Stochastic Volatility Models”. In: *SIAM Journal on Financial Mathematics* 10.2, pp. 491–511. DOI: 10.1137/18M1167565.

- El Euch, Omar and Mathieu Rosenbaum (2019). “The Characteristic Function of Rough Heston Models”. In: *Mathematical Finance* 29.1, pp. 3–38. DOI: 10.1111/mafi.12173.
- Engle, Robert F. (1982). “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation”. In: *Econometrica* 50.4, pp. 987–1007. DOI: 10.2307/1912773.
- Engle, Robert F. and Simone Manganelli (2004). “CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles”. In: *Journal of Business & Economic Statistics* 22.4, pp. 367–381. DOI: 10.1198/073500104000000370.
- Fissler, Tobias and Johanna F. Ziegel (2016). “Higher Order Elicitability and Osband’s Principle”. In: *The Annals of Statistics* 44.4, pp. 1680–1707. DOI: 10.1214/16-AOS1439.
- Forde, Martin and Hongzhong Zhang (2017). “Asymptotics for Rough Stochastic Volatility Models”. In: *SIAM Journal on Financial Mathematics* 8.1, pp. 114–145. DOI: 10.1137/15M1009330.
- Fukasawa, Masaaki (2021). “Volatility Has to Be Rough”. In: *Quantitative Finance* 21.1, pp. 1–8. DOI: 10.1080/14697688.2020.1825781.
- Garcin, Matthieu and Martino Grasselli (2022). “Long versus Short Time Scales: The Rough Dilemma and Beyond”. In: *Decisions in Economics and Finance* 45.1, pp. 257–278. DOI: 10.1007/s10203-021-00358-3.
- Gatheral, Jim, Thibault Jaisson, and Mathieu Rosenbaum (2018). “Volatility Is Rough”. In: *Quantitative Finance* 18.6, pp. 933–949. DOI: 10.1080/14697688.2017.1393551.
- Heston, Steven L. (1993). “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options”. In: *The Review of Financial Studies* 6.2, pp. 327–343. DOI: 10.1093/rfs/6.2.327.
- Mandelbrot, Benoit B. and John W. Van Ness (1968). “Fractional Brownian Motions, Fractional Noises and Applications”. In: *SIAM Review* 10.4, pp. 422–437. DOI: 10.1137/1010093.
- McNeil, Alexander J. and Rüdiger Frey (2000). “Estimation of Tail-Related Risk Measures for Heteroscedastic Financial Time Series: An Extreme Value Approach”. In: *Journal of Empirical Finance* 7.3–4, pp. 271–300. DOI: 10.1016/S0927-5398(00)00012-8.
- Morelli, Giacomo and Paolo Santucci de Magistris (2019). “Volatility Tail Risk under Fractionality”. In: *Journal of Banking & Finance* 108, p. 105654. DOI: 10.1016/j.jbankfin.2019.105654.
- Patton, Andrew J., Johanna F. Ziegel, and Rui Chen (2019). “Dynamic Semiparametric Models for Expected Shortfall (and Value-at-Risk)”. In: *Journal of Econometrics* 211.2, pp. 388–413. DOI: 10.1016/j.jeconom.2018.10.008.
- Pitera, Marcin, Thorsten Schmidt, and Łukasz Stettner (2026). “A Novel Scaling Approach for Unbiased Adjustment of Risk Estimators”. In: *Journal of Empirical Finance* 88, p. 101741. DOI: 10.1016/j.jempfin.2026.101741.
- Rockafellar, R. Tyrrell and Stanislav Uryasev (2000). “Optimization of Conditional Value-at-Risk”. In: *The Journal of Risk* 2.3, pp. 21–41. DOI: 10.21314/JOR.2000.038.
- Sadefo Kamdem, Jules (2005). “Value-at-Risk and Expected Shortfall for Linear Portfolios with Elliptically Distributed Risk Factors”. In: *International Journal of Theoretical and Applied Finance* 8.5, pp. 537–551. DOI: 10.1142/S0219024905003104.

- Sadefo Kamdem, Jules (2012). “VaR and ES for Linear Portfolios with Mixture of Generalized Laplace Distributions Risk Factors”. In: *Annals of Finance* 8.1, pp. 123–150. DOI: 10.1007/s10436-009-0138-6.
- Sadefo Kamdem, Jules and Alan Genz (2008). “Approximation of Multiple Integrals over Hyperboloids with Application to a Quadratic Portfolio with Options”. In: *Computational Statistics & Data Analysis* 52.7, pp. 3389–3407. DOI: 10.1016/j.csda.2007.12.006.
- Tudor, Ciprian A. and Nakahiro Yoshida (2023). “High Order Asymptotic Expansion for Wiener Functionals”. In: *Stochastic Processes and their Applications* 164, pp. 443–492. DOI: 10.1016/j.spa.2023.07.001.