

Roughly Rough: Tractable Markov Processes Emulating Rough Volatility

Gregory Pelts

March 22, 2024

Abstract

We extend conventional Brownian motions by incorporating additional clocks defined as running maxima or, more generally, local times. We speculate that the methodology can emulate rough diffusion. In addition, interpreting the physical time in this way provides a computationally efficient way to introduce jumps. We demonstrate that the approach is capable of handling many stylized facts pertaining to stochastic volatility.

Keywords — rough volatility, fractional Brownian motion, stopping time, jump-diffusion

Introduction

Motivation

The steepening of the at-the-money implied-volatility skew at short option maturities motivates models with more than a single diffusion time scale. Rough-volatility models can produce a skew proportional to $T^{H-1/2}$, where T is the option tenor and $H < 1/2$ is the volatility roughness exponent; this is a statement about the skew, not the level of at-the-money volatility [1, 2, 4]. Fitting SPX options together with VIX options imposes a further constraint: the model must describe both stock returns and the distribution of conditional future variance.

The reciprocal construction below also permits independent smooth and rough diffusion strengths, drift terms and stock–variance correlations. This separates the

leverage contributed by the two clocks, while retaining the simpler mixed-clock construction as an exact subfamily. It is motivated by differing short- and long-expiry skew requirements; calibration, rather than either skew slope alone, must determine the appropriate contributions.

A separate alternative lets a smooth reciprocal volatility model evolve in process, or magic, time and models physical observation time by an affine boundary clock. The resulting stock and volatility jumps may supply short-expiry smile and skew without requiring the same dispersion of conditional future variance as a directly rough volatility clock. This is a hypothesis to assess through SPX calibration and a subsequent variance and VIX-put review; it is not an empirical equivalence between continuous rough volatility and jumps. The two-clock specification and its one-volatility-state conditional variance calculation are given in section 4.1. Its simulation uses adaptive zero-return cycles: the sampled affine endpoint determines a conditional-mean physical-clock increment, and the financial states evolve over the complete intervening magic duration. The accumulated generated physical time is retained across expiries; near-expiry refinement reduces the observation-time discrepancy without resetting the clock to the requested date. The conditioning, numerical quadrature and remaining clock approximation are specified in (4.263)–(4.266).

Fractional and Volterra volatility models provide one response to these observations. Their dependence on past innovations, however, makes simulation and conditional forecasting more demanding than in a finite-dimensional Markov model. The aim here is to retain a tractable Markov state while allowing locally irregular clock activity. This is an alternative construction, not a claim that fractional dynamics are the only explanation of observed scaling. Rough volatility also need not make the stock price a fractional process: the usual construction integrates its volatility against an ordinary Brownian motion, subject to the conditions needed for a true stock-price martingale.

What should be rough?

Volatility level, volatility roughness, and return–volatility dependence are distinct empirical quantities. Realized-volatility studies find variation in estimated roughness across market regimes, and SPX option calibrations associate higher volatility with higher fitted Hurst exponents [3, 4]. These results motivate state-dependent prominence of rough fluctuations; they do not establish a universal rule that rising prices cause greater roughness or that falling prices eliminate it. Calibrated Hurst exponents are model dependent, and a change in their finite-window estimates need not imply a change in a structural local exponent.

The evidence about prices requires a separate distinction. Brandi and Di Mat-

teo [5] find that rough Bergomi can generate price multiscaling, but its relationship with volatility roughness differs from that in their empirical sample. Such finite-scale moment behavior does not by itself identify a singular stock clock. Roughness observed in high-frequency prices can also involve microstructure effects: Chong, Delerue, and Li [6] model rough observation noise alongside an assumed Itô efficient price. That decomposition neither identifies the mechanism proposed here nor establishes that the efficient price must have the same roughness as the observed series.

We therefore retain both calendar-time and rough-clock contributions to the underlying, alongside a rough reciprocal volatility state. In the hybrid specification (4.64), the stock accumulates variance through both dt and $d\mathfrak{t}$, while the reciprocal state ζ has calendar and rough drift but diffusion only in $\mathfrak{t}$. Its rough drift and diffusion strengths can be specified independently, including the zero-diffusion limit. The calendar-only stock model $\mathfrak{V}_S = 0$ remains a nested benchmark, rather than a restriction of the general model. The additional stock channel is a hypothesis to assess against option prices and return dynamics; the empirical studies do not uniquely identify it or impose an ordering of the calendar and rough leverage correlations. A continuous local-martingale construction with a singular clock is also distinct from directly using fractional Brownian motion as a tradable price.

Do we live in volatile time?

What if we make the time itself volatile? Consider a stochastic process

$$d\tau_t = dt + \mathfrak{N}dW_t.$$

Here, dW_t is a standard Wiener process in calendar time t

$$dW_t \sim N\left(0, \sqrt{dt}\right).$$

This process cannot be regarded as a stopping time directly as it is not weakly increasing. Instead, consider its running maximum

$$\tau^{\sup} = \sup_{p < t} \tau_p.$$

and the associated Wiener process

$$dW_{\tau^{\sup}} \sim N\left(0, \sqrt{d\tau^{\sup}}\right).$$

At the point when $\tau_{t_0}^{\sup} = \tau_{t_0}$, the increment of $W_{\tau^{\sup}}$ would scale as

$$\delta W_{\tau^{\sup}} \sim \sqrt{\delta \tau^{\sup}} \sim \delta t^{\frac{1}{4}}.$$

This is reminiscent of the fractional Brownian motion with the Hurst exponent $H = \frac{1}{4}$.

1 Properties of the Running Maximum

Consider the joint distribution of the terminal value of $\tau_{t'}$ and its maximum value $\tau_{[t_0, t']}^{\sup}$ within the time interval $[t_0, t']$. Note that the conditional distribution of $\tau_{t'}$ does not depend on drift. In the driftless case, the mirror image technique produces:

$$\mathbb{P}[\tau_{[t_0, t']}^{\sup} > b | \tau_{t'} = \tau'] = \begin{cases} e^{-2 \frac{(b - \tau')(b - \tau_{t_0})}{\aleph^2(t' - t_0)}} & b > \max(\tau', \tau_{t_0}) \\ 1 & \text{otherwise} \end{cases}.$$

In application to the global running maximum

$$\tau_t^{\sup} \equiv \tau_{[-\infty, t]}^{\sup}$$

this implies

$$\mathbb{P}[\delta\tau^{\sup} > b | \tau_{t'} - \tau_{t_0} = \delta\tau] = \begin{cases} e^{-2 \frac{(b - \delta\tau)b}{\aleph^2\delta t}} & b > \max(\delta\tau, \xi_0) \\ 1 & \text{otherwise} \end{cases}.$$

Here, we introduce the drawdown process $\xi \equiv \tau^{\sup} - \tau$, and denote increments as

$$\delta t = t' - t_0, \quad \delta\tau = \tau_{t'} - \tau_{t_0}, \quad \delta\tau^{\sup} = \tau_{t'}^{\sup} - \tau_{t_0}^{\sup}, \quad \delta\xi = \xi_{t'} - \xi_{t_0}.$$

Increment of the running maximum

The conditional expectation of the increment is given by

$$\begin{aligned} \mathbb{E}[\delta\tau^{\sup} | \delta\tau] &= \int_0^\infty \mathbb{P}[\delta\tau^{\sup} > b | \delta\tau] db \\ &= \begin{cases} \sqrt{\frac{1}{2}\pi\delta t} \aleph e^{\frac{(\delta\tau)^2}{2\aleph^2\delta t}} N\left(-\frac{2\xi_0 - \delta\tau}{\aleph\sqrt{\delta t}}\right) & \delta\tau < \xi_0 \\ \sqrt{\frac{1}{2}\pi\delta t} \aleph e^{\frac{(\delta\tau)^2}{2\delta t\aleph^2}} N\left(-\frac{\delta\tau}{\aleph\sqrt{\delta t}}\right) + \delta\tau - \xi_0 & \text{otherwise} \end{cases}. \end{aligned}$$

Here, $N(*)$ denotes the standard normal CDF. The unconditional expectation of the running maximum increment can be obtained as an expectation of the conditional one:

$$\begin{aligned} \mathbb{E}[\delta\tau^{\sup}] &= \int_{-\infty}^\infty \mathbb{E}\{\delta\tau^{\sup} | \delta\tau\} \frac{e^{-\frac{(\delta\tau - \xi_0)^2}{2\aleph^2\delta t}}}{\aleph\sqrt{2\pi\delta t}} d\tau \\ &= \left(\frac{\aleph^2}{2} + \delta t - \xi_0\right) N\left(\frac{\delta t - \xi_0}{\aleph\sqrt{\delta t}}\right) - \frac{\aleph^2}{2} e^{\frac{2\xi_0}{\aleph^2}} N\left(-\frac{\delta t + \xi_0}{\aleph\sqrt{\delta t}}\right) \\ &\quad + \aleph\sqrt{\frac{\delta t}{2\pi}} e^{-\frac{(\delta t - \xi_0)^2}{2\aleph^2\delta t}}. \end{aligned} \tag{1.1}$$

Drawdown distribution

The process τ can influence the dynamics of observable processes only via its increments; i.e., it is defined modulo a constant shift. The drawdown process ξ is invariant under such shifts and, therefore, can be used to represent the state of the model. The cumulative conditional terminal drawdown distribution is equal to

$$\mathbb{P}[\xi' > \xi | \tau] = \begin{cases} e^{-2\frac{(\xi+\delta\tau)\xi}{\aleph^2\delta t}} & \delta\tau > \xi_0 - \xi \\ 1 & \text{otherwise} \end{cases}.$$

The unconditional cumulative distribution of ξ can be obtained as an expectation of a conditional one:

$$\mathbb{P}[\xi' > \xi] = e^{-\frac{2\xi}{\aleph^2}} N\left(\frac{\delta t - \xi - \xi_0}{\aleph\sqrt{\delta t}}\right) + N\left(-\frac{\delta t + \xi - \xi_0}{\aleph\sqrt{\delta t}}\right), \quad (1.2)$$

or

$$\mathbb{P}[\xi' < \xi] = -e^{-\frac{2\xi}{\aleph^2}} N\left(\frac{\delta t - \xi - \xi_0}{\aleph\sqrt{\delta t}}\right) + N\left(\frac{\delta t + \xi - \xi_0}{\aleph\sqrt{\delta t}}\right). \quad (1.3)$$

Finally, the distribution density can be obtained by differentiation

$$\begin{aligned} \phi_t(\xi) = & \frac{1}{2\aleph^2} e^{-\frac{2\xi}{\aleph^2}} N\left(\frac{\delta t - \xi - \xi_0}{\aleph\sqrt{\delta t}}\right) + \\ & \frac{1}{\aleph\sqrt{2\pi\delta t}} \left(e^{-\frac{(\delta t + \xi + \xi_0)^2 - 4\xi_0\delta t}{2\aleph^2\delta t}} + e^{-\frac{(\delta t + \xi - \xi_0)^2}{2\aleph^2\delta t}} \right). \end{aligned} \quad (1.4)$$

Long-term Limit and Stationary Distribution

The drawdown process is similar to that of the mean reverting processes. In the long-term limit, it approaches a stationary distribution

$$\phi_\infty(\xi) \equiv \lim_{t \rightarrow \infty} \phi_t(\xi) = \frac{2}{\aleph^2} e^{-\frac{2\xi}{\aleph^2}} \quad (1.5)$$

with a mean

$$\lim_{t \rightarrow \infty} \mathbb{E}\{\xi_t\} = \frac{\aleph^2}{2}.$$

The time scale for approaching the stationary state is of the order of $\aleph^2$.

Taking into account the relations between increments in the drawdown and running maximum

$$\delta\xi = \delta\tau^{\text{sup}} - \delta\tau,$$

the expected values of these processes are related as follows

$$\mathbb{E}\{\xi_t\} = \xi_0 + \mathbb{E}\{\delta\tau^{\text{sup}}\} - \delta\tau.$$

2 Quasi-Rough Stochastic Volatility Model

Assuming deterministic interest rates and proportional dividends, the problem can be reduced to one with an underlying martingale process. We shall define it as follows:

$$\begin{aligned} d\tau &= dt - \aleph dW_t^0, \\ dS &= S (\nu dW_{\tau^{\text{sup}}} + \sigma dW_t^1), \\ \langle dW_t^0 W_t^1 \rangle &= \rho dt, \\ \xi &\equiv \tau^{\text{sup}} - \tau. \end{aligned} \tag{2.1}$$

Here, dW_t^i ($i = 0, 1$) denotes a pair of correlated standard Wiener processes, and $dW_{\tau^{\text{sup}}}$ denotes a Wiener process in the running maximum clock. The constant positive model parameter ν is associated with the contribution of the *rough* volatility and σ represents the *smooth* volatility. Without the smooth component, the process would almost always be motionless.

Smooth volatility and correlation

In general, smooth volatility σ and correlation ρ processes can be specified in terms of Brownian motions that run in various clocks. However, in the case of single-factor stochastic volatility, we should assume that they are both functions of the drawdown process. We will assume that the correlation ρ is constant and that the smooth volatility σ is of the following functional form:

$$\sigma = \left(\frac{\gamma}{\aleph^2} \xi \right)^\alpha e^{\frac{\beta}{\aleph^2} \xi}. \tag{2.2}$$

Here, α , β , and γ are constant model parameters in addition to $\aleph$, ν , and ρ . These parameters are scaled in a way that makes them dimensionless. The parameter α is related to the elasticity, which is equal to $1 - \frac{1}{2\alpha}$. For instance, $\alpha = \frac{1}{2}$ is associated with a smooth volatility process, which is normal in the vicinity of zero, $\alpha = 1$ is associated with a Heston-like process, and $\alpha = -1$ with the $\frac{3}{2}$ process. Note that for $\alpha < 0$ this interpretation requires $\beta < 0$. With $\beta > 0$ the volatility process is always asymptotically log-normal for large volatility.

Low volatility regime

The vanilla option market generally implies elasticity below 1. In this case, short-term volatilities appear to spend too much time roaming around zero, contradicting the empirical data. In our model, short-term volatilities do not disappear when ξ

approaches the zero boundary. Instead, in this regime, volatility is dominated by the increasing rough contribution, i.e., the expected increment of τ^{sup} (see Fig. 1).




Fig 1: Rough fair variance

VIX skew

Another problem with standard low elasticity models is predicting a downward VIX skew, while the VIX option market implies the opposite. Let us demonstrate how the proposed form of rough diffusion may address the issue. The rough contribution to the fair variance decreases in ξ . Assuming that the smooth contribution increases with ξ , this prevents VIX from approaching zero with $\xi \rightarrow 0$. Distributions with lower bounds manifest upward skew for strikes that are low enough. The effect is getting stronger in short tenors with larger rough contributions. VIX is based on a relatively short (one month) tenor.

Vol-of-vol term-structure

Another important aspect traditional stochastic vol models miss is the stylized fact that log-normal vol-of-vol scales like T^θ , with $\theta \approx -0.5$. Let us test this behavior by plotting a logarithm of normal vol-of-vol against a logarithm of tenor (see Appendix H). To simplify the analysis, we study the volatilities implied by the fair variances instead of the ATM options. The rough contribution to the fair variance is due to the expected increase in running maximum. To simplify the estimation of the smooth contribution, we will impose $\alpha = 1$, $\beta = 0$, so that $\sigma \sim \xi$. We do not expect that these parameters will have much effect. Then the expectation of a smooth contribution to the fair variance is equal to

$$\text{vol}^2 \sim \mathbb{E} \left\{ \int_0^T \gamma^2 \xi_t dt \right\} = \gamma^2 \int_0^T \mathbb{E} \{ \xi_t \} dt.$$

We observe that, in logarithmic form, the smooth contribution is approximately linear with a negative factor in all tenors. For large ξ , the relation gradually breaks down for tenors below the rough time scale, while it is better preserved for large tenors. However, the low-tenors are expected to dominate by the rough contribution anyway. The latter is approximately linear in all ξ for all tenors.

Simulation

In its parsimonious form, the state of the model is specified by two variables S and ξ . Formally, the simulation requires 3 random numbers per time step: 2 (smooth) + 1 (rough). However, in reality, 2 random numbers suffice. The second Wiener process can be decomposed into a combination of parallel and idiosyncratic portions:

$$dW_t^1 = \rho dW_t^0 + \sqrt{1 - \rho^2} dW_t^\perp \quad \langle dW_t^\perp dW_t^0 \rangle = 0.$$

Consequently, the asset price driver also decomposes as follows

$$dx_t = \rho \sigma dW_t^0 + dx_t^\perp,$$

where

$$dx_t^\perp \equiv \sqrt{1 - \rho^2} dW_t^\perp + \nu dW_{\tau^{\text{sup}}}.$$

The second component can be simulated with a single random number with an increment of quadratic variation is equal to:

$$d \langle x_t^\perp \rangle = (1 - \rho^2) \sigma^2 dt + \nu^2 d\tau^{\text{sup}}.$$

The reference implementation with *pytorch* can be found in Appendix C.

Kolmogorov equation

Let us now derive a backward PDE on option values. Under zero interest rates, in the absence of cash flows, an option value is associated with a martingale process $\psi_t = \psi(t, \xi, S)$. By the standard arguments outside the boundary $\xi = 0$ it must satisfy

$$\left(\partial_t - \partial_\xi + \frac{1}{2} \aleph^2 \partial_\xi^2 + \frac{1}{2} \sigma^2 S^2 \partial_S^2 + \rho \aleph \sigma S \partial_\xi \partial_S \right) \psi = 0 \quad (\xi > 0). \quad (2.3)$$

Here, we assume that smooth volatility is a function of the drawdown process $\sigma = \sigma(\xi)$, as in Eq. 2.2. The rough volatility contribution manifests itself via the following Neumann boundary condition:

$$\left(\partial_\xi + \frac{1}{2} \nu^2 S^2 \partial_S^2 \right) \psi \Big|_{\xi=0} = 0. \quad (2.4)$$

Extensions

State-dependent rough diffusion driver

The process Eq. 2.1 can be generalized as

$$\begin{aligned} d\tau &= \mu dt - \aleph dW_t^0, \\ dS &= S (dW_{\tau^{\text{sup}}} + \sigma dW_t^1), \\ \langle dW_t^0 dW_t^1 \rangle &= \rho dt, \\ \xi &\equiv \tau^{\text{sup}} - \tau, \end{aligned} \quad (2.5)$$

with μ , $\aleph$, σ , all being arbitrary positive value functions of ξ or even separate processes. In addition, the correlation ρ can also be a function of ξ or form a separate process. Then the equations (2.3), (2.4) should be modified as follows

$$\begin{aligned} \left(\partial_t - \mu \partial_\xi + \frac{1}{2} \aleph^2 \partial_\xi^2 + \frac{1}{2} \sigma^2 S^2 \partial_S^2 + \rho \aleph \sigma S \partial_\xi \partial_S \right) \psi &= 0 \quad (\xi > 0), \\ \left(\partial_\xi + \frac{1}{2} S^2 \partial_S^2 \right) \psi \Big|_{\xi=0} &= 0. \end{aligned}$$

Note that we no longer need to introduce the parameter ν , as it can be absorbed by the other parameters and rescaling of the process ξ . The model can be mapped to one with constant rough volatility $\aleph = 1$ by rescaling the physical time as follows

$$dt = \aleph^2 d\tilde{t}.$$

Indeed, under the rescaled time, the τ -process becomes

$$d\tau = \mu \aleph^{-2} d\tilde{t} - dW_{\tilde{t}}.$$

To keep the process perpetually rough, we should impose

$$\lim_{\xi \rightarrow 0} \frac{\mu}{\aleph^2} < \infty.$$

Otherwise, eventually, the drawdown process would not be able to return to zero, causing the running maximum process to stop. Furthermore, the reciprocal $\aleph^2$ should be integrable with the ξ distribution for the model with constant parameters with $\aleph = 1$ defined in Eq. 1.4:

$$\int \phi_t \aleph^{-2} d\xi < \infty.$$

Otherwise, the physical time would reach infinity at the final value of the rescaled time $\tilde{t}$. This would imply that in physical time, the rescaled time process $\tilde{t}$ freezes, eliminating rough dynamics. It is sufficient to test this condition with stationary distribution in Eq. 1.5:

$$\int e^{-2\xi} \aleph^{-2} d\xi < \infty.$$

If divergence occurs at $\xi \rightarrow 0$, the τ process will freeze after reaching the running maximum. If a divergence occurs at $\xi \rightarrow \infty$ the τ -process would never be able to return to running maximum. Assume that in the vicinity of $\xi = 0$

$$\lim_{\xi \rightarrow 0} \aleph \sim \xi^\lambda.$$

Here, λ is allowed to be negative. However, to avoid divergence, we should impose $\lambda < 1$.

System of nested clocks

The dynamic can be extended with additional clocks defined via running maxima of other processes. These processes may already have rough components. This would create a hierarchy of clocks that may be used to emulate fractional diffusion with various Hurst exponents representable as negative powers of two or a sum of those. Later in the paper, we will demonstrate a more parsimonious way to handle this issue.

Stochastic physical time

The physical time itself can be modeled as one of the stochastic clocks defined in terms of some formal *model* time, as proposed in [7]. This would provide an organic and computationally efficient way to introduce jumps.

Stochastic local vol

A standard local stochastic vol extension can be applied to capture the vanilla option market accurately:

$$\nu = \nu(t, S) = \phi(t, S)\tilde{\nu}, \quad \sigma = \sigma(\xi, t, S) = \phi(t, S)\tilde{\sigma}(\xi).$$

Here, ϕ is a leverage function, and $\tilde{\nu}, \tilde{\sigma}$ are borrowed from the pure stochastic vol process.

3 Quasi-Rough Diffusion: Generic Framework

Edge time and local martingale

Let $\mathfrak{L}_t$ be a continuous positive weakly growing local martingale process, i.e., such that

$$\mathbb{E}[\mathfrak{L}_T] \geq \mathfrak{L}_t.$$

Such processes can be uniquely decomposed into a sum of a global martingale $\mathfrak{G}_T$ and a weakly increasing process $\mathfrak{T}_t$.

$$\mathfrak{L}_t = \mathfrak{G}_t + \mathfrak{T}_t. \quad (3.1)$$

To demonstrate this, let us investigate the process in the accrued quadratic variation (variance) clock $\langle \mathfrak{L} \rangle$. Then, due to the local martingale property, everywhere outside the boundary $\mathfrak{L}_t = 0$ the process must be of the form

$$d\mathfrak{L} = dW_{\langle \mathfrak{L} \rangle}. \quad (3.2)$$

Here, $W_{\langle \mathfrak{L} \rangle}$ denotes the standard Wiener process in the variance clock. The solution of the equation is

$$\mathfrak{L} = \left| \tilde{W}_{\langle \mathfrak{L} \rangle} \right|, \quad (3.3)$$

where $\tilde{W}_{\langle \mathfrak{L} \rangle}$ is another standard Wiener related to the original one as follows

$$dW_{\langle \mathfrak{L} \rangle} = \text{sgn} \left(\tilde{W}_{\langle \mathfrak{L} \rangle} \right) d\tilde{W}_{\langle \mathfrak{L} \rangle}.$$

According to the Tanaka formula,

$$\mathfrak{L} = W_{\langle \mathfrak{L} \rangle} + L \left[\tilde{W} = 0 \right]_{\langle \mathfrak{L} \rangle} = W_{\langle \mathfrak{L} \rangle} + \frac{1}{2} L [\mathfrak{L} = 0]_{\langle \mathfrak{L} \rangle}. \quad (3.4)$$

Here, $L[*]$ denotes the local time spent in the vicinity of the event. Thus we found a solution to the decomposition Eq. 3.1 with

$$\mathfrak{r}_t = \frac{1}{2}L[\mathfrak{L} = 0]_{\langle \mathfrak{L} \rangle}.$$

We will prefer another equivalent representation of the almost always positive solution to Eq. 3.2 alternative to Eq. 3.2

$$\mathfrak{L} = W_{\langle \mathfrak{L} \rangle} - W_{\langle \mathfrak{L} \rangle}^{\text{inf}},$$

as it is more in line with our previously proposed *rough* volatility driver associated with a drawdown process. Assume that $\mathfrak{L}$ is a time-homogeneous Markov process without additional state variables. Then its volatility is a function of $\mathfrak{L}$ through a stationary relation

$$d\langle \mathfrak{L} \rangle = \sigma^2(\mathfrak{L})dt,$$

the process $\mathfrak{r}_t$ can be generally associated with appropriately scaled local time spent at the proximity of the boundary $\mathfrak{L} = 0$. We will call this an edge clock.

Physical time process in the edge clock

Consider physical time as an increasing process as a stochastic process $t_{\mathfrak{r}}$ in the edge clock. Assume that $\mathfrak{L}_t$ is a one-dimensional time-homogeneous Markov process in the physical clock. The only state variable $\mathfrak{L}$ is always zero when the edge clock moves forward, making it invisible from the perspective of an observer living in the edge clock. This makes $t_{\mathfrak{r}}$ a process with independent stationary increments, that is, a Levy process. This is also a pure jump process, as it contains an infinite number of jumps within any edge time interval. As with any Levy process, it can be fully specified by its Laplace exponent representation:

$$\mathbb{E} \left[e^{\lambda \delta t} \right] = e^{\Psi(\lambda) \delta \mathfrak{r}} \quad (\lambda \in \mathbb{C}). \quad (3.5)$$

For an increasing process, $\Psi(\lambda)$ is an analytic function whose domain of definition is contained in the complex half-plane $\Re \lambda < 0$. We approximate this process with a discrete-time process defined as follows. Each time step of the discrete process consists of two consecutive parts. The first part starts at the boundary and ends after the physical time interval equal to T . The second part is completed when the process returns to the boundary. The total increase in physical time in one step is equal to

$$\delta t = T + T',$$

where T' is the time elapsed before returning to the boundary. According to Eq. 3.1, the expected increase in edge time in the first part of the step is equal to the expectation of $\mathfrak{L}$ at the end of the first part of the step:

$$\mathbb{E}[\mathfrak{r}_1 - \mathfrak{r}_0] = \mathbb{E}[\mathfrak{L}_1].$$

At the second part of the step, the process $\mathfrak{L}$ remains positive and, therefore, the process $\mathfrak{r}$ remains constant. Thus, the expected total growth of the edged time in a single step is equal to

$$\mathbb{E}[\delta \mathfrak{r}] = \mathbb{E}[\mathfrak{L}_1].$$

The process $t_{\mathfrak{L}}$ is a limit $T \rightarrow 0$ of this discrete-time process with a step frequency equal to $\mathbb{E}[\mathfrak{L}_1]^{-1}$. Therefore, for its Laplace exponent representation Eq. 3.5 we should have:

$$\Psi(\lambda) = \lim_{T \rightarrow 0} \mathbb{E}[\mathfrak{L}_1]^{-1} \left(\log \mathbb{E} \left[e^{\lambda T'} \right] + \lambda T \right). \quad (3.6)$$

We will show that, generally, in the short-term limit, the time taken to return to the boundary is an order of magnitude larger than the time spent leaving it $T' \gg T$. This reduces the expression Eq. 3.6 to

$$\Psi(\lambda) = \lim_{T \rightarrow 0} \mathbb{E}[\mathfrak{L}_1]^{-1} \log \mathbb{E} \left[e^{\lambda T'} \right] = \lim_{\xi \rightarrow 0} \mathfrak{L}_1(\xi)^{-1} \log \psi_{\xi}(\lambda). \quad (3.7)$$

Here $\psi_{\xi}(\lambda)$ denotes Laplace transformation of the time elapsed before reaching the boundary from the initial state ξ :

$$\psi_{\xi}(\lambda) \equiv \mathbb{E} \left[e^{\lambda T'} \middle| \xi_1 = \xi \right]. \quad (3.8)$$

Reverse direction: edge time in the physical clock

By the duality $\{t_{\mathfrak{r}} > t\} = \{\mathfrak{r}_t < \mathfrak{r}\}$, the distribution of the edge time $\mathfrak{r}_t$ starting from the boundary $\xi_0 = 0$ is uniquely determined by the Laplace exponent Ψ of Eq. 3.5. We denote its moment and cumulant generating functions from the boundary by

$$M_{\mathfrak{r}_t}(\lambda) \equiv \mathbb{E} \left[e^{\lambda \mathfrak{r}_t} \middle| \xi_0 = 0 \right], \quad K_{\mathfrak{r}_t}(\lambda) \equiv \log M_{\mathfrak{r}_t}(\lambda). \quad (3.9)$$

More generally, for any $\mathbb{R}_+$ -valued random variable X we write f_X , F_X , M_X , and K_X for its density, cumulative distribution function, MGF, and CGF, respectively. Taking a Laplace transform of Eq. 3.9 in t yields the double transform

$$\int_0^{\infty} e^{\Lambda t} M_{\mathfrak{r}_t}(\lambda) dt = -\frac{\Psi(\Lambda)}{\Lambda(\Psi(\Lambda) + \lambda)}, \quad \Re \Lambda, \Re \lambda \leq 0. \quad (3.10)$$

Thus Eq. 3.5 encodes all cumulants of physical time in the edge clock as Taylor coefficients of Ψ at zero whenever Ψ is analytic there, while Eq. 3.10 encodes all cumulants of edge time in the physical clock (from the boundary) through the same Ψ .

Singular change of numeraire

Consider the generic continuous diffusion process on $\mathbb{R}_+$ of the form

$$d\xi_t = \mu(\xi)dt + \sigma(\xi)dW_t.$$

Any local martingale must satisfy the backward Kolmogorov equation

$$\partial_t \psi + \mathfrak{D}\psi = 0, \quad (3.11)$$

where we define the elliptic differential operator

$$\mathfrak{D} = \mu \partial_\xi + \frac{1}{2} \sigma^2 \partial_\xi^2,$$

which can be decomposed as

$$\mathfrak{D} = \mathcal{D}\partial, \quad \mathcal{D} = \mu \cdot + \frac{1}{2} \sigma^2 \partial_\xi. \quad (3.12)$$

Here the symbol $\cdot$ indicates a multiplication operator. The distribution density ϕ satisfies the adjoint forward Fokker-Planck equation

$$\partial_t \phi - \mathfrak{D}^* \phi = 0, \quad (3.13)$$

where

$$\mathfrak{D}^* = -\partial \mathcal{D}^*, \quad \mathcal{D}^* = \mu \cdot - \frac{1}{2} \partial_\xi \sigma^2. \quad (3.14)$$

The value of the stationary local martingale must satisfy

$$\mathfrak{D}\psi = 0.$$

We will allow the boundary to be attainable. We also consider the process to be fully reflecting (nonabsorbing). This prevents $\mathfrak{L}$ from being a global martingale. The martingale property would preclude the process from leaving the boundary once it is reached, which would make the process fully absorbing. We will study nonstationary solutions of this equation with the following asymptotic behavior near the boundary.

$$\psi = {}^0\psi + {}^1\psi \mathfrak{L} + o(\mathfrak{L}). \quad (3.15)$$

Here, ${}^0\psi$ and ${}^1\psi$ are some constants. They can be obtained as

$${}^0\psi = \lim_{\xi \rightarrow 0} \psi, \quad {}^1\psi = \lim_{\xi \rightarrow 0} \frac{\psi - {}^0\psi}{\mathfrak{L}}. \quad (3.16)$$

Solutions associated with global martingales must satisfy ${}^1\psi = 0$. Although $\mathfrak{L}$ is not a global martingale, it can still act as a Radon–Nikodym derivative.

However, the measure produced by such transformation would not be equivalent to the original. Rather, as we will demonstrate later, it can be associated with a certain conditional process. The formal application of the Girsanov theorem results in the following modification of the drift term:

$$\mu \rightarrow \tilde{\mu} = \mu + \mathbf{q},$$

where

$$\mathbf{q} = \sigma^2 \partial_\xi \log \mathfrak{L}. \quad (3.17)$$

Indeed, the formulae Eq. 3.11, Eq. 3.13 are equivalent to

$$\partial_t \tilde{\psi} = -\tilde{\mathfrak{D}} \tilde{\psi} = -\tilde{\mathcal{D}} \partial \tilde{\psi}, \quad \partial_t \tilde{\phi} = \tilde{\mathfrak{D}}^* \tilde{\phi} = -\partial \tilde{\mathcal{D}} \tilde{\psi}, \quad (3.18)$$

where

$$\psi = \tilde{\psi}, \quad \tilde{\mathcal{D}} = \tilde{\mu} \cdot + \frac{1}{2} \sigma^2 \cdot \partial_\xi = \tilde{\mathcal{D}} + \mathbf{q} \cdot, \quad (3.19)$$

$$\tilde{\phi} = \mathfrak{L} \phi, \quad \tilde{\mathcal{D}}^* = \tilde{\mu} \cdot - \frac{1}{2} \partial_\xi \sigma^2 \cdot = \mathcal{D}^* + \mathbf{q} \cdot, \quad (3.20)$$

as long as $\mathfrak{L}$ solves to the Kolmogorov equation, i.e., $\mathfrak{D}\mathfrak{L} = 0$. We will call the process associated with the above singular measure change a dual process. The stationary solution to the adjoint Fokker-Planck equation Eq. 3.13 must be of the form

$$\mathcal{D}^* \phi = \text{const} \equiv -{}_0\phi.$$

The same analysis applied in the dual representation yields

$$\tilde{\mathcal{D}}^* \tilde{\phi} = \text{const} \equiv {}_1\phi.$$

Applying here Eq. 3.20, we obtain

$${}_1\phi = (\mathcal{D}^* + \mathbf{q}) (\mathfrak{L} \phi) = \mathfrak{L} (\mathcal{D}^* + \mathbf{q}) \phi - \frac{\sigma^2}{2} \partial_\xi \mathfrak{L} \partial_\xi \phi = \mathfrak{L} (-{}_0\phi + \frac{\mathbf{q}}{2} \phi).$$

This leads to the following representation for the generic stationary solution of the forward equation:

$$\phi = \frac{2}{\mathbf{q}} \left(\frac{{}_1\phi}{\mathfrak{L}} + {}_0\phi \right). \quad (3.21)$$

For nonstationary solutions, this decomposition still applies in the form of an asymptotic expansion near the boundary.

$$\phi(t, \xi)|_{\xi \rightarrow 0} = \frac{2}{\mathbf{q}} \left(\frac{{}_1\phi(t)}{\mathfrak{L}} + {}_0\phi(t) + o(1) \right) \quad (3.22)$$

as long as the initial condition satisfies the constraint. The expansion factors can be determined as

$${}_1\phi = \frac{1}{2} \lim_{\xi \rightarrow 0} \mathfrak{q} \mathfrak{L} \phi, \quad {}_0\phi = \frac{1}{2} \lim_{\xi \rightarrow 0} \mathfrak{q} \left(\phi - \frac{{}_0\phi}{\mathfrak{L}} \right). \quad (3.23)$$

Assuming that this asymptotic relation is differentiable and homogeneous in time, its application in conjunction with Eq. 3.14 yields

$$\partial_\xi \mathcal{D}^* \phi|_{\xi \rightarrow 0} = \mathcal{O} \left(\frac{2}{\mathfrak{q} \mathfrak{L}} \right).$$

Integrating this relation over ξ we obtain

$$\mathcal{D}^* \phi|_{\xi \rightarrow 0} = -{}_0\phi + \mathcal{O}(\varepsilon) = -{}_0\phi + o(1). \quad (3.24)$$

Here we have introduced the function

$$\varepsilon(\xi') \equiv \int_0^{\xi'} \frac{2}{\mathfrak{q} \mathfrak{L}} d\xi. \quad (3.25)$$

Similar analysis applied to the dual process yields

$$\tilde{\mathcal{D}}^* \tilde{\phi}|_{\xi \rightarrow 0} = {}_1\phi + o(\tilde{\varepsilon}), \quad (3.26)$$

where

$$\tilde{\varepsilon}(\xi') \equiv \int_0^{\xi'} \frac{2}{\mathfrak{q}} d\xi. \quad (3.27)$$

Let ϕ represent the distribution density for the process with nonabsorbing boundary. Then

$$\int_0^\infty \phi(T, \xi) d\xi = 1$$

Differentiation of the right part of this with respect to time while applying Eq. 3.13 together with Eq. 3.24 results in

$$\frac{d}{dT} \int_0^\infty \phi(T, \xi) d\xi = \int_0^\infty -\partial \mathcal{D}^* \phi d\xi = \lim_{\xi \rightarrow 0} \mathcal{D}^* \phi = -{}_0\phi. \quad (3.28)$$

Thus we have obtained the following boundary condition for forward equation associated with the non-absorbing (fully reflecting) process

$${}_0\phi = 0. \quad (3.29)$$

Performing an analogous analysis in the dual representation, we arrive at

$$\frac{d}{dt} \int_0^\infty \phi \mathfrak{L} d\xi = \frac{d}{dt} \int_0^\infty \tilde{\phi} d\xi = \int_0^\infty -\partial \mathcal{D}^* \tilde{\phi} d\xi = \lim_{\xi \rightarrow 0} \mathcal{D}^* \tilde{\phi} = {}_1\phi. \quad (3.30)$$

This result can be related to an expectation of the future value of the local martingale $\mathfrak{L}$. In addition, with an opposite sign, it can be interpreted as the absorption rate. Note that according to Eq. 3.22 for a positive distribution ϕ it is required that

$${}_1\phi \geq 0. \quad (3.31)$$

This makes absorption for the dual process impossible or, in other words, makes the boundary unattainable. In fact, the dual process can be interpreted as the base process conditional on never reaching the boundary. Consider the boundary condition ${}_1\phi = 0$. According to Eq. 3.22, it is equivalent to

$$\phi|_{\xi \rightarrow 0} = \mathcal{O}\left(\frac{2}{q}\right),$$

which means that it is an order of magnitude smaller than the density for an absorbing process. This boundary condition can be associated with a complete absorption at the boundary. The probability of not touching the boundary can be calculated via this solution as follows

$$\mathbb{P}[\forall_{t < T} : \xi_t > 0] = \int_0^\infty \phi(T, \xi) d\xi.$$

Differentiating this equality with respect to time while applying Eq. 3.28 produces the following formula for the first touch distribution density

$$-\frac{d}{dT} \mathbb{P}[\forall_{t < T} : \xi_t > 0] = -{}_0\phi.$$

Relation to local time

According to our definition of edge time $\mathfrak{t}$ in Eq. 3.1, its expected growth rate is equal to the expected growth rate of the local martingale $\mathfrak{L}$, which, in turn, can be calculated with Eq. 3.30

$$\frac{d}{dT} \mathbb{E}[\mathfrak{t}_T] = \frac{d}{dT} \mathbb{E}[\mathfrak{L}_T] = {}_1\phi(T) \quad (3.32)$$

Here $\phi(T, \xi)$ is the distribution density for the base process. Application of the L'Hôpital's rule to the ratio of the cumulative distribution function to the ε function Eq. 3.25 yields

$$\lim_{\xi \rightarrow 0} \frac{\mathbb{P}[\xi_T < \xi]}{\varepsilon(\xi)} = \lim_{\xi \rightarrow 0} \frac{\phi(T, \xi)}{\frac{2}{q(\xi)\mathfrak{L}(\xi)}} = \frac{1}{2} \lim_{\xi \rightarrow 0} \phi(T, \xi) q(\xi) \mathfrak{L}(\xi) = {}_1\phi(T).$$

Here we also applied Eq. 3.23. Integrating Eq. 3.32 over time after substituting there the above equality we obtain

$$\mathbb{E}[\mathfrak{r}_t] = \mathbb{E}\left[\lim_{\xi \rightarrow 0} \frac{1}{\varepsilon(\xi)} \int \mathcal{I}\{\xi_t < \xi\} dt\right].$$

Here, $\mathcal{I}\{*\}$ denotes an indicator function. Note that the variable under expectation on the left side of the equation is continuous and weakly increasing, thus falling under our axiomatic definition of the edge time. This leads to the conclusion:

$$\mathfrak{r}_t = \lim_{\xi \rightarrow 0} \frac{1}{\varepsilon(\xi)} \int \mathcal{I}\{\xi_t < \xi\} dt. \quad (3.33)$$

Boundary conditions

We will say that a process is quasi-rough if its drift and/or quadratic variation have edge time contribution:

$$dx_t = \mu_x dt + \mathfrak{I} d\mathfrak{r} + \sigma_x dW_t + \aleph dW_\tau.$$

Application of the Ito lemma to a derived process $\psi_t = \psi(t, \xi_t, x)$ produces

$$\begin{aligned} d\psi = & \left(\partial_t + \mu \partial_\xi + \mu_x \partial_x + \frac{\sigma^2}{2} \partial_\xi^2 + \frac{\sigma_x^2}{2} \partial_x^2 + \rho \sigma \sigma_x \partial_\xi \partial_x \right) \psi dt \\ & + \left(\mathfrak{I} \partial_x + \frac{\aleph^2}{2} \partial_x^2 \right) \psi d\mathfrak{r} + \sigma_\xi \partial_\xi \psi dW_t + \sigma_x \partial_x \psi dW_t^x + \aleph \frac{\partial \psi}{\partial x} dW_\tau. \end{aligned}$$

Here ρ denotes the correlation between the processes x and ξ . As the edge time does not move outside the boundary $\xi = 0$ the standard equation on local martingale processes still applies:

$$\left(\partial_t + \mu \partial_\xi + \mu_x \partial_x + \frac{\sigma^2}{2} \partial_\xi^2 + \frac{\sigma_x^2}{2} \partial_x^2 + \rho \sigma \sigma_x \partial_\xi \partial_x \right) \psi = 0. \quad (3.34)$$

Here we assume that the process ξ is Markov with no extra state variables. The result obtained here can be trivially extended to higher-dimensional models. Edge time can only contribute via the boundary condition. To derive this condition, let us apply the asymptotic decomposition of ψ near the boundary, generalizing the analogous decomposition Eq. 3.15:

$$\psi(t, \xi, x)_{\xi \rightarrow 0} = {}^0\psi(t) + {}^1\psi(t)\mathfrak{L} + o(\mathfrak{L}).$$

with the decomposition coefficients defined in Eq. 3.16 The boundary condition is obtained by replacing $d\mathfrak{r}$ with $d\mathfrak{L}$ in the limit $\xi \rightarrow 0$:

$$\lim_{\xi \rightarrow 0} \left(\mathfrak{I} \partial_x \psi + \frac{\aleph^2}{2} \partial_x^2 \psi \right) = {}^1\psi. \quad (3.35)$$

Quasi-rough stochastic rates

Edge time can also be used as a contributor to accrual for the money market account defined as follows

$$\Phi_{t'} = \Phi_{t_0} \exp \left(\int_{t_0}^{t'} r(t) dt + \mathfrak{r} \delta \mathfrak{t} \right). \quad (3.36)$$

Here $r(t)$ is the conventional instantaneous risk-free rate and $\mathfrak{r}$ is the intensity of quasi-rough contribution. This methodology allows one to model large and volatile long-term rates while keeping the short-term instantaneous rate at a low or zero level. This feature can be particularly helpful in modeling a low-rate regime. The rough contribution does not affect the Kolmogorov equation, which remains the same as for the standard diffusion:

$$\left(\partial_t + \mu \partial_\xi + \frac{\sigma^2}{2} \partial_\xi^2 - r \cdot \right) \psi = 0. \quad (3.37)$$

However, now it must be paired with the following boundary condition:

$${}^1\psi = \mathfrak{r} {}^0\psi, \quad (3.38)$$

or, equivalently,

$$\psi(t, \xi)|_{\xi \rightarrow 0} = {}^0\psi (1 + \mathfrak{r} \mathfrak{L} + o(\mathfrak{L})).$$

The adjoint boundary condition for the associated forward equation should be defined in such a way that if satisfied, the following condition must hold:

$$\int_0^\infty \psi \phi d\xi = \text{const}$$

Differentiating the left side of the equation under the assumption that the function under the integral satisfies the backward and forward equations Eq. 3.11, Eq. 3.13 produces

$$\begin{aligned} \frac{d}{dt} \int_0^\infty \psi \phi d\xi &= \int_0^\infty \{-\phi \mathfrak{D} \psi + \psi \mathfrak{D}^* \phi(T, \xi)\} d\xi \\ &= - \int_0^\infty \partial_\xi \left\{ \mu \phi \psi + \frac{1}{2} (\phi \sigma^2 \partial_\xi \psi - \psi \partial_\xi (\phi \sigma^2)) \right\} d\xi \\ &= \lim_{\xi \rightarrow 0} \left\{ \mu \phi \psi + \frac{1}{2} (\phi \sigma^2 \partial_\xi \psi - \psi \partial_\xi (\phi \sigma^2)) \right\} \\ &= \lim_{\xi \rightarrow 0} \left(\psi \mathcal{D}^* \phi + \frac{1}{2} \phi \sigma^2 \partial_\xi \psi \right). \end{aligned}$$

Alternatively, in the dual representation, this expression can be represented as

$$\frac{d}{dt} \int_0^\infty \psi \phi d\xi = \frac{d}{dt} \int_0^\infty \tilde{\psi} \tilde{\phi} d\xi = \lim_{\xi \rightarrow 0} \left(\tilde{\psi} \tilde{\mathcal{D}}^* \tilde{\phi} + \frac{1}{2} \tilde{\phi} \sigma^2 \partial_\xi \tilde{\psi} \right).$$

This representation is better equipped to handle the contribution of the terms with the multiplier $\mathfrak{L}$ in ψ as it is associated with a constant function in the dual representation. In the absence of the constant term in ψ the contribution would be

$$\lim_{\xi \rightarrow 0} \left(\tilde{\psi} \tilde{\mathcal{D}}^* \tilde{\phi} + \frac{1}{2} \tilde{\phi} \sigma^2 \partial_\xi \tilde{\psi} \right) = {}^1\psi \lim_{\xi \rightarrow 0} \tilde{\mathcal{D}}^* \tilde{\phi} = {}^1\psi {}_1\phi.$$

Here we also applied Eq. 3.26. The constant term is easier to handle in the base representation. In the absence of the rough terms in ψ the contribution would be

$$\lim_{\xi \rightarrow 0} \left(\psi \mathcal{D}^* \phi + \frac{1}{2} \phi \sigma^2 \partial_\xi \psi \right) = {}^0\psi \lim_{\xi \rightarrow 0} \mathcal{D}^* \phi = -{}^0\psi {}_0\phi.$$

Here we also applied Eq. 3.24. Finally, combining it all together, we obtained the following results.

$$\frac{d}{dt} \int_0^\infty \psi \phi d\xi = {}^1\psi {}_1\phi - {}^0\psi {}_0\phi. \quad (3.39)$$

In order to eliminate the boundary contribution, we should impose on the forward equation a boundary condition orthogonal to Eq. 3.38 with respect to the inner product on the left side of Eq. 3.39, which is:

$${}_0\phi = \mathfrak{I} {}_1\phi. \quad (3.40)$$

In the limit $\mathfrak{I} \rightarrow \infty$ the boundary conditions for the Kolmogorov and Fokker-Planck equations become

$${}^0\psi = 0, \quad {}_1\phi = 0. \quad (3.41)$$

According to Eq. 3.16, Eq. 3.23 this is equivalent to

$$\lim_{\xi \rightarrow 0} \psi = 0, \quad \lim_{\xi \rightarrow 0} \mathfrak{q} \mathfrak{L} \phi = 0. \quad (3.42)$$

As expected, these conditions are associated with full absorption. In credit applications, the methodology generalizes the Merton capital structure model [8]. In the general case $\mathfrak{I} < \infty$ absorption is partial. That is, the process can reach the boundary without necessarily triggering default. The extreme case $\mathfrak{I} = \infty$ is associated with the conventional version of the Merton model.

Example: Running maximum

Let us apply the methodology to the previously studied running maximum process. Outside the boundary, the associated drawdown process follows the following stochastic differential equation.

$$d\xi = -dt + \aleph dW. \quad (3.43)$$

In application to this process, the operator Eq. 3.12 is equal to

$$\mathcal{D} = 1 \cdot + \aleph \partial_\xi.$$

The local martingale solution to this equation that turns zero at the boundary is unique modulo the a multiplier, which we choose to make the associated edge time equal to the running maximum:

$$\mathfrak{L} = \frac{\aleph^2}{2} \left(e^{\frac{2\xi}{\aleph^2}} - 1 \right). \quad (3.44)$$

An application of the generic framework Eq. 3.17Eq. 3.25 yields

$$\mathfrak{q} = \frac{2}{1 - e^{-\frac{2\xi}{\aleph^2}}}, \quad \varepsilon = 1 - e^{-\frac{2\xi}{\aleph^2}} = \frac{2\xi}{\aleph^2} + o(\xi)$$

Therefore, according to Eq. 3.33, the edge-time is equal to the local time spent at the boundary multiplied by the factor $\frac{\aleph^2}{2}$. This result is consistent with the Tanaka formula for a running maximum.

Let us calculate the Laplace transform of the ξ_1 distribution

$$\Phi_T(\Lambda) \equiv \mathbb{E} \left[e^{\Lambda \xi_1} \right]. \quad (3.45)$$

Consequently, performing integration by parts under the assumption $\Re \Lambda < 0$, we obtain

$$\begin{aligned} \Phi_T(\Lambda) &= \int_0^\infty e^{\Lambda \xi} d\mathbb{P} [\xi_1 < \xi] = -\frac{1}{\Lambda} \int_0^\infty e^{\Lambda \xi} \mathbb{P} [\xi_1 < \xi] d\xi = \\ &= \frac{2}{2 - \aleph^2 \Lambda} \left((1 - \aleph^2 \Lambda) N \left(\frac{(\aleph^2 \Lambda - 1) \sqrt{T}}{\aleph} \right) e^{T\Lambda(\frac{1}{2}\aleph^2 \Lambda - 1)} + N \left(\frac{\sqrt{T}}{\aleph} \right) \right) \end{aligned} \quad (3.46)$$

The expected increase of edge time is equal to

$$\begin{aligned} \mathbb{E} [\delta \mathfrak{r}_1] &= \mathbb{E} [\mathfrak{L}_1] = \frac{\aleph^2}{2} \left(\lim_{\Lambda \rightarrow \frac{2}{\aleph^2}} \Phi_T(\Lambda) - 1 \right) \\ &= \left(\frac{\aleph^2}{2} + T \right) N \left(\frac{\sqrt{T}}{\aleph} \right) - \frac{\aleph^2}{2} N \left(-\frac{\sqrt{T}}{\aleph} \right) + \aleph \sqrt{\frac{T}{2\pi}} e^{-\frac{T}{2\aleph^2}}. \end{aligned} \quad (3.47)$$

Here we use the definition of $\mathfrak{L}$ process in Eq. 3.44. This value is exactly the same as in Eq. 1.1. This confirms that the edge time associated with the local martingale Eq. 3.44 is indeed the same as the running maximum. In the short-term limit the Laplace transformation becomes

$$\Phi_T(\Lambda)_{T \rightarrow 0} = 1 + \aleph \sqrt{\frac{2T}{\pi}} \Lambda + O(T). \quad (3.48)$$

This implies that in the short-term limit the distribution of ξ_1 is approaching deterministic with the value

$$\xi_1 \approx \aleph \sqrt{\frac{2T}{\pi}}.$$

The short term limit of Eq. 3.47 yields

$$\mathbb{E}[\delta \mathfrak{r}_1]_{T \rightarrow 0} = \aleph \sqrt{\frac{2T}{\pi}} + O(T). \quad (3.49)$$

This result can be also obtained by differentiating Eq. 3.48 with respect to Λ and using the following relation

$$\mathbb{E}[\delta \mathfrak{r}_1] = \mathbb{E}[\delta \mathfrak{L}_1] = \left. \frac{d}{d\Lambda} \Phi_T \right|_{\Lambda=0}.$$

The Laplace transform of the distribution of the return time is equal to

$$\mathbb{E}[e^{\lambda T'}] = \int_0^\infty \psi(\xi, \lambda) d\mathbb{P}[\xi_1 < \xi], \quad (3.50)$$

where $\psi(\xi, \lambda)$ is defined in Eq. 3.8. Differentiating both sides w.r.t. the step starting time and applying the Ito lemma, we arrive at

$$\left(-\partial_\xi + \frac{1}{2} \aleph^2 \partial_\xi^2 + \lambda \right) \psi = 0.$$

Solving this differential equation with the boundary condition $\psi(0, \lambda) = 1$ we obtain

$$\psi(\xi, \lambda) = e^{\xi \frac{1 - \sqrt{1 - 2\aleph^2 \lambda}}{\aleph^2}}. \quad (3.51)$$

This confirms the well-known result for the Laplace transform for the inverse Gaussian distribution. Substituting Eq. 3.51 to Eq. 3.50 and applying notations Eq. 3.45 we find that

$$\mathbb{E}[e^{\lambda T'}] = \int_0^\infty e^{\xi \frac{1 - \sqrt{1 - 2\aleph^2 \lambda}}{\aleph^2}} d\mathbb{P}[\xi_1 < \xi] = \Phi_T \left(\frac{1 - \sqrt{1 - 2\aleph^2 \lambda}}{\aleph^2} \right). \quad (3.52)$$

Applying here the asymptotic expansion Eq. 3.48 we observe

$$\lim_{T \rightarrow 0} \mathbb{E} \left[e^{\lambda T'} \right] = 1 + \sqrt{\frac{2T}{\pi}} \frac{1 - \sqrt{1 - 2\aleph^2 \lambda}}{\aleph^2} + O(T). \quad (3.53)$$

Counterintuitively, in the short-time limit, the time taken to return to the boundary is an order of magnitude larger than the time spent leaving it. Finally, substituting Eq. 3.52, Eq. 3.49 and Eq. 3.48 into Eq. 3.6 we obtain the Laplace exponent representation of the physical time process in the edge clock:

$$\begin{aligned} \Psi(\lambda) &= \lim_{T \rightarrow 0} \left(\aleph \sqrt{\frac{2T}{\pi}} \right)^{-1} \left(\log \Phi_T \frac{1 - \sqrt{1 - 2\aleph^2 \lambda}}{\aleph^2} + \lambda T \right) \\ &= \lim_{T \rightarrow 0} \left(\aleph \sqrt{\frac{2T}{\pi}} \right)^{-1} \left(\aleph \sqrt{\frac{2T}{\pi}} \frac{1 - \sqrt{1 - 2\aleph^2 \lambda}}{\aleph^2} + \lambda T \right) = \frac{1 - \sqrt{1 - 2\aleph^2 \lambda}}{\aleph^2}. \end{aligned}$$

This agrees with the well-known result for the inverse Gaussian process.

Distribution of the edge time from the boundary

Since $\Psi(\lambda) = (1 - \sqrt{1 - 2\aleph^2 \lambda})/\aleph^2$ is the classical inverse-Gaussian Levy exponent, the physical time $t_{\mathfrak{r}}$ admits the concrete Brownian representation

$$t_{\mathfrak{r}} \stackrel{d}{=} \inf\{u \geq 0 : u + \aleph W_u = \mathfrak{r}\},$$

where W is a standard Brownian motion. Inverting a first-passage process yields the running maximum, so

$$\mathfrak{r}_t \stackrel{d}{=} \sup_{u \leq t} (u + \aleph W_u),$$

consistent with the identification $\mathfrak{r} = \tau^{\sup}$ of the auxiliary process of Eq. 2.1. The corresponding CDF at $\xi_0 = 0$ is given by Eq. 1.3, and its density by Eq. 1.4. This is the running-maximum analogue of the affine Mittag-Leffler identification of Eq. 4.29 below: in both cases $\mathfrak{r}_t$ is the inverse of the physical-time subordinator; the Mittag-Leffler law arises only in the stable (affine) case, whereas here the inverse of an inverse-Gaussian Levy process gives the running-maximum-of-Brownian-drift distribution instead.

Since $\Psi(\lambda)$ is analytic at $\lambda = 0$, all cumulants of physical time in the edge clock are finite. Expanding,

$$\Psi(\lambda) = \lambda + \frac{\aleph^2}{2} \lambda^2 + \frac{\aleph^4}{2} \lambda^3 + O(\lambda^4),$$

and reading off the Taylor coefficients from Eq. 3.5,

$$\frac{\mathbb{E}[t_{\mathfrak{r}}]}{\delta\mathfrak{r}} = 1, \quad \frac{\text{Var}[t_{\mathfrak{r}}]}{\delta\mathfrak{r}} = \mathfrak{N}^2, \quad \frac{\kappa_3[t_{\mathfrak{r}}]}{\delta\mathfrak{r}} = 3\mathfrak{N}^4.$$

The reverse direction follows from Eq. 3.10: differentiating $M_{\mathfrak{r}_t}$ with respect to λ at $\lambda = 0$ and inverting the Laplace transform in t produces

$$\mathbb{E}[\mathfrak{r}_t | \xi_0 = 0] = t + \frac{\mathfrak{N}^2}{2} + o(1), \quad t \rightarrow \infty,$$

which agrees with the large-tenor limit of Eq. 3.47. In the short-tenor limit, the drift is negligible and $\mathfrak{r}_t \stackrel{d}{=} \mathfrak{N}\sqrt{t}|Z|$ with $Z \sim \mathcal{N}(0,1)$, giving $\mathbb{E}[\mathfrak{r}_t | \xi_0 = 0] = \mathfrak{N}\sqrt{2t/\pi} + O(t)$, consistent with Eq. 3.49.

4 Quasi-Rough Extension of Affine Process

Motivation

Earlier, we have demonstrated how to generate a process that emulates rough diffusion of the Hurst index equal to $\frac{1}{2}$. However, empirical evidence suggests that volatility often manifests a more significant roughness associated with much smaller Hurst indices. This issue can be addressed by layering running-maximum-linked clocks. In this way, we can emulate rough dynamics with Hurst indices equal to a sum of negative powers of two. This would require additional state variables with a total number equal to the maximum negative power, thus increasing complexity and adding calibration instability. Fortunately, replacing running maximum with a more general edge-time technique provides a much more parsimonious solution that does not require any extra state variable. Furthermore, as we shall see, an edge time obtained from a one-dimensional affine process allows one to emulate behavior associated with an arbitrary Hurst index. Affine models, such as the Heston stochastic volatility model and Cox-Ingersoll-Ross stochastic rate models, are very well studied. They offer an array of closed form formulae for such important quantities as forward variances and forward rates. This feature is essential to model VIX options and interest rate derivatives. Although these formulae are not directly applicable in the presence of rough components, they can still be used to assess the smooth contribution. As we will demonstrate, various quantities pertaining to the edge time, and, therefore, linked to the rough contribution to forward variances and rates, are also available in closed form or can be efficiently approximated.

Mean-reversion free process

Let us start by studying an affine process with no mean reversion:

$$d\xi = \theta dt + \sqrt{2\xi} dW_t. \quad (4.1)$$

Here θ is a positive constant. We have reduced the volatility parameter by appropriately scaling the underlying variable. As we will show later, most of the results obtained in mean reversion free environments can be extended to a generic case. The local martingale solutions for this process satisfy the following Kolmogorov equation:

$$\partial_t \psi + \theta \partial_\xi \psi + \xi \partial_\xi^2 \psi = 0. \quad (4.2)$$

For our purposes, the boundary must be attainable, which requires the Feller condition to be violated, that is,

$$\theta < 1. \quad (4.3)$$

Then, the Kolmogorov equation has a unique up-to-a-scale solution turning zero at the boundary:

$$\mathfrak{L} = \xi^{1-\theta}. \quad (4.4)$$

Applying the generic formulae Eq. 3.17, Eq. 3.25, Eq. 3.27 with this solution we obtain

$$\mathfrak{q}(\xi) = 2(1 - \theta), \quad \varepsilon(\xi) = \frac{\xi^\theta}{\theta(1 - \theta)}, \quad \tilde{\varepsilon}(\xi) = \frac{\xi}{1 - \theta}. \quad (4.5)$$

Dimensional analysis and scale symmetry

Perform a dimensional analysis of Eq. 4.1. Based on the random walk term $(\xi)^2 \sim (t)(\xi)$ and therefore $(\xi) \sim (t)$, i.e., ξ is of the same dimension as time. Then, based on the drift term, the parameter θ is dimensionless. This is consistent with the way it participates in the Feller condition Eq. 4.3. According to our definition of the edge time, its dimension is the same as the dimension of the local stationary martingale solution Eq. 4.4. Therefore,

$$\delta\mathfrak{r} \sim \mathfrak{L} \sim \xi^{1-\theta} \sim (\delta t)^{1-\theta}.$$

Consequently, the model is invariant under the following scaling symmetry:

$$\delta t \rightarrow \beta \delta t, \quad \xi \rightarrow \beta \xi, \quad \delta\mathfrak{r} \rightarrow \beta^{1-\theta} \delta\mathfrak{r} \quad (\beta > 0). \quad (4.6)$$

This scaling matches the self-similarity $B_{\beta t}^H \stackrel{d}{=} \beta^H B_t^H$ of fractional Brownian motion of Hurst index $H = (1 - \theta)/2$: the rough Wiener increment $\delta W_{\mathfrak{r}} \sim (\delta\mathfrak{r})^{1/2} \sim (\delta t)^{(1-\theta)/2}$ transforms with the same exponent as B^H itself, providing structural support for the hypothesized link between the affine construction and fractional diffusion of that Hurst index across the entire family $\theta \in (0, 1)$.

Distribution of the time to reach the boundary

Let $T' = t_2 - t_1$ denote the time elapsed before returning to the boundary from the initial state $\xi = \xi_1$. According to the scaling symmetry, the cumulative distribution function for this variable must be of the form

$$\mathbb{P}[T' < T | \xi_1 = \xi] = \Omega\left(\frac{\xi}{T}\right). \quad (4.7)$$

Note that, for the fixed maturity t_2 , the cumulative probability is a martingale and, therefore, is subject to the Kolmogorov equation Eq. 4.2. In application to the right part of Eq. 4.7 it leads to

$$(z + \theta)\Omega'(z) + z\Omega''(z) = 0. \quad (4.8)$$

Solving this equation with the boundary conditions

$$\Omega(0) = 1, \quad \lim_{z \rightarrow \infty} \Omega(z) = 0$$

we obtain:

$$\Omega(z) = \frac{\Gamma(1 - \theta, z)}{\Gamma(1 - \theta)}. \quad (4.9)$$

Here, $\Gamma(*)$ and $\Gamma(*, *)$ are the gamma function and the upper incomplete gamma function, respectively. This means that the variable $z_1 \equiv \frac{\xi_1}{T'}$ follows the gamma distribution with a shape parameter equal to $1 - \theta$. Now we are ready to calculate the Laplace transform of the return-time distribution:

$$\begin{aligned} \psi_\xi(\lambda) &= - \int_0^\infty e^{\lambda \frac{\xi}{z}} d\Omega(z) = \frac{1}{\Gamma(1 - \theta)} \int_0^\infty z^{-\theta} e^{\lambda \frac{\xi}{z} - z} d\Omega(z) \\ &= \frac{2(-\lambda\xi)^{\frac{1-\theta}{2}}}{\Gamma(1 - \theta)} K_{\theta-1}\left(2\sqrt{-\lambda\xi}\right). \end{aligned} \quad (4.10)$$

Here K_α denotes the modified Bessel function of the second kind of the order α . In the limit $\xi \rightarrow 0$ this result is approximated by

$$\psi_\xi(\lambda)_{\xi \rightarrow 0} = 1 + \frac{\Gamma(-1 + \theta)}{\Gamma(1 - \theta)} (-\lambda\xi)^{1-\theta} - \frac{1}{\theta} \lambda\xi + O\left(\xi^{2-\theta}\right). \quad (4.11)$$

The physical time process in the edge clock

According to our generic conclusion, in the edge clock, the physical time follows the increasing pure jump Levy process with the Laplace exponent representation Eq. 3.7. After substituting there Eq. 4.4 and Eq. 4.11 we obtain

$$\Psi(\lambda) = \lim_{\xi \rightarrow 0} \xi^{\theta-1} \log \psi_\xi(\lambda) = \frac{\Gamma(-1 + \theta)}{\Gamma(1 - \theta)} (-\lambda)^{1-\theta}. \quad (4.12)$$

Note, that the cumulant generator $\Psi(\lambda)$ branches at $\lambda = 0$. This implies that all the cumulants of the Levy process are infinite. This is caused by the lack of mean reversion.

Green function

Green function $G(t', \xi' | t_0, \xi_0)$ is the distribution of $\xi_{t'} = \xi'$ subject to the initial condition $\xi_{t_0} = \xi_0$. The Green function satisfies the Kolmogorov equation Eq. 4.2 with respect to the second pair of arguments. With respect to the first pair of arguments, it satisfies the adjoint Fokker-Planck equation:

$$(-\partial_t - \theta \partial_\xi + \partial_\xi^2 \xi \cdot) \phi = 0. \quad (4.13)$$

The Green function for this process is related to the noncentral χ^2 distribution

$$G(t', \xi' | t_0, \xi_0) = \frac{1}{t' - t_0} \left(\frac{\xi'}{\xi_0} \right)^{\frac{\theta-1}{2}} e^{-\frac{\xi' + \xi_0}{t' - t_0}} I_{\theta-1} \left(2\sqrt{\frac{\xi \xi_0}{T^2}} \right). \quad (4.14)$$

Note the following special cases

$$\lim_{\xi_0 \rightarrow 0} G(t', \xi' | t_0, \xi_0) = \frac{1}{\Gamma(\theta)} e^{-\frac{\xi'}{t' - t_0}} (t' - t_0)^{-\theta} \xi'^{\theta-1} \quad (4.15)$$

$$G(t', \xi' | t_0, \xi_0)_{\xi' \rightarrow 0} = \frac{1}{\Gamma(\theta)} e^{-\frac{\xi_0}{t' - t_0}} (t' - t_0)^{-\theta} \xi'^{\theta-1} + \mathcal{O}(\xi'^\theta). \quad (4.16)$$

These results can be also obtained by leveraging the scaling symmetry Eq. 4.6, thus reducing the problem to solving ordinary differential equations. In terms of the generic asymptotic decomposition for solutions of the Kolmogorov and Fokker-Planck equations Eq. 3.16, Eq. 3.22 they can be written as follows:

$$\begin{aligned} {}^0G(t', \xi' | t_0) &= \lim_{\xi_0 \rightarrow 0} G(t', \xi' | t_0, \xi_0) \\ &= \frac{1}{\Gamma(\theta)} e^{-\frac{\xi'}{t' - t_0}} (t' - t_0)^{-\theta} \xi'^{\theta-1}, \end{aligned} \quad (4.17)$$

$$\begin{aligned} {}_1G(t' | t_0, \xi_0) &= \frac{1}{2} \lim_{\xi' \rightarrow 0} \mathfrak{q}(\xi') \mathfrak{L}(\xi') G(t', \xi' | t_0, \xi_0) \\ &= -\frac{1}{\Gamma(\theta - 1)} e^{-\frac{\xi_0}{t' - t_0}} (t' - t_0)^{-\theta}. \end{aligned} \quad (4.18)$$

Here, we applied Eq. 4.5. The generic formula Eq. 3.32 for the expected increase in edge time, in terms of the Green function, can be formulated as follows.

$$\frac{d}{dt} \mathbb{E}[\mathfrak{r}_t | \xi_{t_0} = \xi_0] = {}_1G(t | t_0, \xi_0), \quad (4.19)$$

which in this case yields

$$\frac{d}{dt} \mathbb{E}[\mathfrak{r}_t | \xi_{t_0} = \mathbf{x}_0] = -\frac{1}{\Gamma(\theta - 1)} e^{-\frac{\xi_0}{t' - t_0}} (t' - t_0)^{-\theta}. \quad (4.20)$$

Scale invariant parametrization

The Green function Eq. 4.14 can be rewritten in the following manifestly scale-invariant form

$$G(t', \xi' | t_0, \xi_0) = \frac{\mathcal{G}(Z(\xi', t'; t_0), Z(\xi_0, t_0; t'))}{Y(t', t_0)} \quad (4.21)$$

where

$$Y(t', t_0) \equiv |t' - t_0|, \quad Z(t', \xi'; t_0) \equiv \frac{\xi'}{Y(t', t_0)}, \quad (4.22)$$

$$\mathcal{G}(z', z_0) \equiv \left(\frac{z'}{z_0}\right)^{\frac{\theta-1}{2}} e^{-z_0 - z'} I_{\theta-1}(2\sqrt{z' z_0}). \quad (4.23)$$

The edge cases Eq. 4.15, Eq. 4.16 can be derived from the asymptotic expansion of the Bessel function

$$I_{\theta-1}(2\sqrt{z}) = z^{\frac{\theta-1}{2}} \left(\frac{1}{\Gamma(\theta)} + \mathcal{O}(z) \right).$$

In the scale-invariant form they can be formulated as

$${}_0G(t', \xi' | t_0) = \frac{e^{-Z(\xi', t'; t_0)} Z(\xi', t'; t_0)^{\theta-1}}{\Gamma(\theta) Y(t', t_0)}, \quad (4.24)$$

$${}_1G(t' | t_0, \xi_0) = \frac{e^{-Z(\xi_0, t_0; t')}}{-\Gamma(\theta - 1) Y(t', t_0)^\theta}. \quad (4.25)$$

The previously obtained result for the distribution of time to reach the boundary Eq. 4.9, Eq. 4.9 also naturally formulates in the new notations:

$$\mathbb{P}[\{\xi_t = 0\} \succ \{t = t'\} | \xi_{t_0} = \xi_0] = \Omega(Z(t_0, x_0; t')). \quad (4.26)$$

Here, $\{*\}$ denotes the first occurrence of the listed condition, and the symbol $\succ$ specifies the order of the occurrences.

Expected increase in edge time

To model VIX options, we need the ability to handle forward variances directly. In our approach, the rough contribution to the expected variance is proportional to the expected value of the increase in edge time. The expectation of the increment of the edge time is equal to

$$\mathbb{E}[\delta \mathfrak{r} | \xi_{t_0} = \xi_0] = \mathbb{E}[\mathfrak{L}(\xi') | \xi_{t_0} = \xi_0] - \mathfrak{L}(\xi_0) = \int_0^\infty G(t', \xi | t_0, \xi_0) \xi^{1-\theta} d\xi - \xi_0^{1-\theta}.$$

Applying here Eq. 4.15 in the case of $\xi_0 = 0$ yields

$$\mathbb{E}[\delta \mathfrak{r} | \xi_{t_0} = 0] = \int_0^\infty \frac{\xi^{\theta-1} (t' - t_0)^{-\theta} e^{-\frac{\xi}{t' - t_0}}}{\Gamma(\theta)} \xi^{1-\theta} d\xi = \frac{(t' - t_0)^{1-\theta}}{\Gamma(\theta)}.$$

Thus we established that, in this case, the rough component of the variance scales as $(\delta t)^{1-\theta}$. This emulates rough diffusion with a Hurst index $H = \frac{1-\theta}{2}$. The expected increase in edge time can be calculated without resorting to integration over the noncentral χ^2 distribution. The edge time remains constant until the process reaches the boundary, at which point its expected value can be evaluated with the central chi-square distribution Eq. 4.15. As has been demonstrated, the reciprocal time to reach the boundary $z = \frac{\xi_0}{t - t_0}$ follows the gamma distribution Eq. 4.9. Therefore,

$$\begin{aligned} \mathbb{E}[\delta \mathfrak{r} | \xi_{t_0} = \xi_0] &= \int_{\frac{\xi_0}{t-t_0}}^\infty \frac{e^{-z} z^{-\theta}}{\Gamma(1-\theta)} \frac{\left(t - t_0 - \frac{\xi_0}{z}\right)^{1-\theta}}{\Gamma(\theta)} dz \\ &= \frac{\xi_0^{1-\theta}}{\Gamma(1-\theta)\Gamma(\theta)} \int_1^\infty e^{-\frac{t-t_0}{\xi_0} y} (y-1)^{1-\theta} \frac{dy}{y} \\ &= -\xi_0^{1-\theta} \frac{\Gamma(\theta-1, \frac{\xi_0}{t-t_0})}{\Gamma(\theta-1)} \end{aligned} \quad (4.27)$$

The expected growth rate of the edge time can be obtained by differentiating the above result over the expiration time:

$$\frac{d}{dt} \mathbb{E}[\mathfrak{r}_t | \xi_{t_0} = \xi_0] = \frac{(t - t_0)^{-\theta} e^{-\frac{\xi_0}{t-t_0}}}{-\Gamma(\theta-1)}.$$

This is in agreement with Eq. 4.20.

Distribution of the edge time from the boundary

Under the initial condition $\xi_{t_0} = 0$, the scaling symmetry Eq. 4.6 applied to a path started at the boundary implies that the ratio

$$R \equiv \frac{\delta \mathfrak{r}}{(\delta t)^{1-\theta}} \quad (4.28)$$

has a distribution independent of δt . Its law can be identified in closed form as follows.

According to Eq. 3.5 and Eq. 4.12, in the edge clock, the physical time is a stable subordinator of index $1-\theta$. The branching of Ψ at $\lambda = 0$ makes every cumulant of $t_{\mathfrak{r}}$ infinite (contrast with the running maximum case, where all cumulants are finite). The reverse direction is regular: substituting $\Psi(\Lambda) = \Gamma(\theta - 1)/\Gamma(1 - \theta) \cdot (-\Lambda)^{1-\theta}$ into Eq. 3.10 produces the double Laplace transform of $\mathfrak{r}_t$ from the boundary in closed form, and inverting in t shows that $\mathfrak{r}_t$ is the *inverse* of a $(1 - \theta)$ -stable subordinator. Thus R inherits the Mittag-Leffler law of index $1 - \theta$ [11, 10, 12, 9], rescaled by the intensity of $\Psi(\lambda)$:

$$R \stackrel{d}{=} \frac{\Gamma(2 - \theta)}{\Gamma(\theta)} Y, \quad M_Y(\lambda) = E_{1-\theta}(\lambda), \quad (4.29)$$

where E_α denotes the classical Mittag-Leffler function [10]. The Mittag-Leffler law arises here specifically because Ψ is stable; for a general Ψ the inverse subordinator has a different distribution (e.g., the running-maximum-of-Brownian-drift distribution in the running maximum case of Eq. 2.1), and only the double transform Eq. 3.10 continues to apply. The first moment

$$\mathbb{E}[R] = \frac{1}{\Gamma(\theta)}$$

is consistent with the $\xi_0 = 0$ specialization of Eq. 4.27. Indeed, setting $z_0 = \xi_0/\delta t$ and expanding the upper incomplete gamma function for small z_0 ,

$$\Gamma(\theta - 1, z_0) = \Gamma(\theta - 1) - \frac{z_0^{\theta-1}}{\theta - 1} + O(z_0^\theta),$$

substituting into Eq. 4.27 and using $(\theta - 1)\Gamma(\theta - 1) = \Gamma(\theta)$ yields

$$\mathbb{E}[\delta \mathfrak{r} | \xi_0] = -\xi_0^{1-\theta} + \frac{\delta t^{1-\theta}}{\Gamma(\theta)} + o(1).$$

Since $1 - \theta > 0$, $\xi_0^{1-\theta} \rightarrow 0$ as $\xi_0 \rightarrow 0$, and Eq. 4.28 produces $\mathbb{E}[R] = 1/\Gamma(\theta)$, as claimed.

The density of R is not elementary in general, but it can be evaluated from that of the one-sided $(1 - \theta)$ -stable law [13] via the identity $Y \stackrel{d}{=} S^{-(1-\theta)}$, where S has $M_S(\lambda) = e^{-(-\lambda)^{1-\theta}}$ for $\Re \lambda \leq 0$.

As a sanity check, consider the special case $\theta = \frac{1}{2}$ (Hurst index $H = \frac{1}{4}$, matching the running maximum model of Eq. 2.1). Then $\Gamma(2 - \theta)/\Gamma(\theta) = \frac{1}{2}$, and $g_{1/2}$ is the density of the Lévy distribution. A short calculation gives

$$f_R(r) = \frac{2}{\sqrt{\pi}} e^{-r^2}, \quad r > 0,$$

i.e., a half-normal, reproducing the short-time limit Eq. 3.49 of the running maximum increment.

Support for the Hurst-index emulation at arbitrary θ

The Mittag-Leffler identification Eq. 4.29 carries stronger consequences than the mean-scaling argument sketched earlier. Since the Mittag-Leffler moments $\mathbb{E}[Y^n] = n!/\Gamma(1 + n(1 - \theta))$ are finite for every n ,

$$\mathbb{E}[\mathfrak{r}_t^n | \xi_0 = 0] \propto (\delta t)^{n(1-\theta)}, \quad n \geq 1.$$

The Wiener process $W_{\mathfrak{r}_t}$ that drives the rough component of Eq. 2.1 is standard Brownian in the edge clock, so

$$\mathbb{E}[(W_{\mathfrak{r}_t})^n | \xi_0 = 0] \propto (\delta t)^{n(1-\theta)/2} = (\delta t)^{nH}, \quad H = \frac{1 - \theta}{2},$$

which matches the corresponding fractional Brownian motion moment exponents $\mathbb{E}[(B_t^H)^n] \propto t^{nH}$ to all orders. This extends the $n = 2$ variance-scaling argument used earlier to justify the emulation of fractional diffusion of Hurst index $H = (1 - \theta)/2$ to a full-moment consistency statement, for arbitrary $\theta \in (0, 1)$. The scaling symmetry Eq. 4.6 is itself a process-level statement of the same self-similarity that fractional Brownian motion enjoys under $t \rightarrow \beta t$, $B_t^H \rightarrow \beta^H B_t^H$, providing further structural support for the emulation across the entire family.

The extension is a marginal-scaling statement, however, and not a process-level equivalence: the affine construction is Markov in ξ by design, whereas fractional Brownian motion is not. The autocovariance structures therefore differ, and the emulation should be read as reproducing the same self-similar marginal moment scaling as fractional Brownian motion while remaining a tractable Markov process.

Conditional process and duality without mean reversion

Let us study the process under the condition that the boundary was reached at the time t_∞ ¹. This can be accomplished by taking an Arrow-Debreu security linked to the event as a numeraire. In this case, the Radon-Nikodym derivative is the distribution density for the first touch time. According to Eq. 4.9, it is, up to a multiplier, equal to

$$\omega_{t_\infty}(t, \xi) = \Omega'(z) \frac{dz}{dt} = -\frac{1}{\Gamma(1-\theta)} \xi^{1-\theta} (t_\infty - t)^{\theta-2} e^{-z} \quad \left(z = \frac{\xi}{t_\infty - t} \right).$$

Application of the Girsanov theorem to this case produces the following drift adjustment

$$\begin{aligned} \mathbb{E} \left[d\xi \middle| \inf_{\xi_t=0} t = t_\infty \right] &= \mathbb{E}[d\xi] + d\langle \xi, \log \omega_{t_\infty} \rangle = (\theta + 2\xi \partial_\xi \log \omega_{t_\infty}) dt \\ &= \left(2 - \theta - \frac{2\xi}{t_\infty - t} \right) dt. \end{aligned}$$

Note that the conditional process is very similar to the original, except with the modified Feller parameter

$$\theta \rightarrow \tilde{\theta} = 2 - \theta \quad (4.30)$$

and a time-dependent mean reversion. The latter disappears in the limit $t_\infty \rightarrow \infty$, which is associated with the condition of never reaching the boundary

$$\mathbb{E}[d\xi | \forall_t \xi_t > 0] = \tilde{\theta} dt.$$

Note that the modified parameter satisfies the Feller condition $\tilde{\theta} > 1$, making the boundary unattainable as expected. In the general case $t_\infty \neq \infty$, the time-dependent mean reversion can be eliminated with the following change of variables

$$(t, \xi) \rightarrow (\tilde{t}, \tilde{\xi}) : \quad \tilde{t} = \frac{1}{t_\infty - t}, \quad \tilde{\xi} = \xi \frac{d\tilde{t}}{dt} = \frac{\xi}{(t_\infty - t)^2}. \quad (4.31)$$

Then, indeed

$$\begin{aligned} \mathbb{E} \left[d\tilde{\xi} \middle| \inf_{\xi_t=0} t = t_\infty \right] &= \frac{\mathbb{E}[d\xi | \inf_{\xi_t=0} t = t_\infty]}{(t_\infty - t)^2} + \frac{2\xi dt}{(t_\infty - t)^3} = \frac{\tilde{\theta} dt}{(t_\infty - t)^2} = \tilde{\theta} d\tilde{t} \\ d\langle \tilde{\xi} \rangle &= \left(\frac{d\tilde{t}}{dt} \right)^2 d\langle \xi \rangle = \frac{2\xi dt}{(t_\infty - t)^4} = 2\tilde{\xi} d\tilde{t}. \end{aligned} \quad (4.32)$$

In the scale-invariant form, the transformation can be written as

$$(t, \xi) \rightarrow (\tilde{t}, \tilde{\xi}) : \quad \tilde{t} = \frac{1}{|Y(t_\infty, t)|}, \quad \tilde{\xi} = \frac{Z(t, \xi; t_\infty)}{|Y(t_\infty, t)|}. \quad (4.33)$$

¹The reasoning behind these notations will transpire later

Extension to the joint (ξ, S) dynamics. The same Girsanov mechanism produces the conditional drift of any process correlated with ξ . For an asset S driven jointly with ξ through correlated Wiener drivers,

$$d\xi = \theta dt + \sqrt{2\xi} dW_t^\xi, \quad \frac{dS}{S} = \sqrt{\frac{\mathcal{H}}{2}} \xi \left(\rho dW_t^\xi + \sqrt{1 - \rho^2} dW_t^\perp \right), \quad (4.34)$$

under $\mathbb{P}$ we have $d\langle \xi, \log S \rangle = \rho\sqrt{\mathcal{H}}\xi dt$. The martingale part of $d\log \omega_{t_\infty}$ is

$$[d\log \omega_{t_\infty}]_{\text{mart}} = \sqrt{\frac{2}{\xi}} \left[(1 - \theta) - \frac{\xi}{t_\infty - t} \right] dW_t^\xi, \quad (4.35)$$

obtained by adding the martingale parts of $(1 - \theta)\log \xi$ and $-\xi/(t_\infty - t)$. The cross-variation with $\log S$ produces the extra drift acquired by $\log S$ under the conditional measure:

$$\mathbb{E} \left[d\log S \mid \inf_{\xi_t=0} t = t_\infty \right] - \mathbb{E}[d\log S] = \rho\sqrt{\mathcal{H}} \left[(1 - \theta) - \frac{\xi}{t_\infty - t} \right] dt. \quad (4.36)$$

In the never-touch limit,

$$\mathbb{E}[d\log S \mid \forall_t \xi_t > 0] - \mathbb{E}[d\log S] = \rho\sqrt{\mathcal{H}}(1 - \theta) dt, \quad (4.37)$$

i.e. a constant drift on $\log S$. The ξ -dependence in $\sqrt{\mathcal{H}\xi/2} \cdot \sqrt{2/\xi}$ cancels exactly: the shift in $\log S$ carries unit scaling weight and is therefore a scalar under the affine scaling symmetry $(t, \xi) \rightarrow (\beta t, \beta \xi)$, matching the scaling weight of dt .

Written in the scale-invariant coordinates Eq. 4.33, using $\xi/(t_\infty - t) = \tilde{\xi}/\tilde{t}$ and $dt = d\tilde{t}/\tilde{t}^2$ together with the invariance $\tilde{\mathcal{H}}\tilde{\xi}d\tilde{t} = \mathcal{H}\xi dt$ (which fixes $\tilde{\mathcal{H}} = \mathcal{H}/\tilde{t}^4$ and $\sqrt{\tilde{\mathcal{H}}} = \sqrt{\mathcal{H}}/\tilde{t}^2$), the drift takes the manifestly scalar form

$$\mathbb{E} \left[d\log S \mid \inf_{\xi_t=0} t = t_\infty \right] - \mathbb{E}[d\log S] = \rho\sqrt{\tilde{\mathcal{H}}} \left[(1 - \theta) - \frac{\tilde{\xi}}{\tilde{t}} \right] d\tilde{t}. \quad (4.38)$$

This is the same expression as Eq. 4.36 in the dual variables: substituting $\sqrt{\tilde{\mathcal{H}}} = \sqrt{\mathcal{H}}/\tilde{t}^2$ and $d\tilde{t} = \tilde{t}^2 dt$ recovers Eq. 4.36 exactly. At a fixed terminal touch, $t \rightarrow t_\infty$ maps to $\tilde{t} \rightarrow \infty$, but the coefficient $\sqrt{\tilde{\mathcal{H}}}$ then vanishes as $\tilde{t}^{-2}$; this is a coordinate stretching of a finite physical-time interval, not an explosion of the conditional asset process. The never-touch limit is instead $t_\infty \rightarrow \infty$ at fixed t , which maps to $\tilde{t} \rightarrow 0$ and is most directly represented by Eq. 4.37.

The perpendicular driver $W^\perp$ is independent of W^ξ , so its distribution is unaffected by the change of measure. The quadratic variation of $\log S$ is likewise

unaffected: the Girsanov transform of $\mathbb{P} \rightarrow \mathbb{P}^h$ only shifts the drift on W^ξ . Combining, the conditional dynamics of $\log S$ under the never-touch measure take the closed form

$$d \log S = -\frac{1}{2} \frac{\mathcal{H}\xi}{2} dt + \rho \sqrt{\mathcal{H}} (1 - \theta) dt + \sqrt{\frac{\mathcal{H}\xi}{2}} (\rho d\tilde{W}_t^\xi + \sqrt{1 - \rho^2} dW_t^\perp), \quad (4.39)$$

where $\tilde{W}^\xi = W^\xi - \int (1 - \theta) \sqrt{2/\xi_s} ds$ is a standard Brownian motion under the conditional measure. The finite- t_∞ case is obtained by retaining the state-dependent term in Eq. 4.36; it is not a deterministic contribution along an individual conditioned path.

Mean-reversion

The generic one-dimensional affine process

$$d\xi = (\theta - \kappa\xi) dt + \sqrt{2\xi} dW_t. \quad (4.40)$$

can be reduced to the mean-reversion-free versions Eq. 4.1 via the following change of variables

$$(t, \xi) \rightarrow (\tilde{t}, \tilde{\xi}) : \quad \tilde{t} = \frac{e^{\kappa t}}{\kappa}, \quad \tilde{\xi} = \xi \frac{d\tilde{t}}{dt} = e^{\kappa t} \xi. \quad (4.41)$$

Indeed, in new variables, the growth of quadratic variation and expected state variable are equal to

$$\begin{aligned} \mathbb{E} [d\tilde{\xi}] &= e^{\kappa t} (\mathbb{E} [d\xi] + \kappa\xi dt) = \theta e^{\kappa t} dt = \theta d\tilde{t}, \\ d\langle \tilde{\xi} \rangle &= e^{2\kappa t} d\langle \xi \rangle = 2e^{2\kappa t} \xi dt = 2\tilde{\xi} d\tilde{t}. \end{aligned} \quad (4.42)$$

As a consequence, the results Eq. 4.23, Eq. 4.24, Eq. 4.25 and Eq. 4.26 still apply after the mapping in Eq. 4.22 is modified as follows:

$$Z(t', \xi'; t_0) = \frac{\tilde{\xi}'}{|\tilde{t}' - \tilde{t}_0|} = \frac{e^{\kappa t'} \kappa \xi'}{|e^{\kappa t'} - e^{\kappa t}|} = \frac{\kappa \xi'}{|1 - e^{\kappa(t_0 - t')}|} = \frac{\xi'}{|Y(t', t_0)|}, \quad (4.43)$$

where

$$Y(t', t_0) = \frac{e^{\kappa(t_0 - t')} - 1}{\kappa}. \quad (4.44)$$

In these notations, the transformation to the mean-reversion-free parameterization Eq. 4.41 becomes

$$(t, \xi) \rightarrow (Y(t_\infty, t), Y(t_\infty, t) Z(t, \xi, t_\infty)).$$

The anchor point t_∞ can be chosen anywhere in the past or future.

Extension to the joint (ξ, S) system. The change of variables Eq. 4.41 extends to the joint dynamics of ξ and $\log S$ in the presence of mean reversion. Under $(t, \xi) \rightarrow (\tilde{t}, \tilde{\xi})$ the correlated Wiener drivers transform as scalars with weight $\frac{1}{2}$ in $d\tilde{t}$: $dW_t^\xi = e^{-\kappa t/2} d\tilde{W}_t^\xi$ and $dW_t^\perp = e^{-\kappa t/2} d\tilde{W}_t^\perp$, both preserving the correlation ρ . Substituting into Eq. 4.34 together with $\xi dt = \tilde{\xi} d\tilde{t}$ gives

$$d \log S = -\frac{1}{2} \frac{\tilde{\mathcal{H}}}{2} \tilde{\xi} d\tilde{t} + \sqrt{\frac{\tilde{\mathcal{H}}}{2}} \tilde{\xi} (\rho d\tilde{W}_t^\xi + \sqrt{1-\rho^2} d\tilde{W}_t^\perp), \quad \tilde{\mathcal{H}} = \mathcal{H} e^{-2\kappa t}. \quad (4.45)$$

That is, $\log S$ retains its classic-form Heston dynamics in the mean-reversion-free variables with the coefficient $\mathcal{H}$ rescaled by $e^{-2\kappa t}$, exactly as required by scale invariance: under Eq. 4.41 the pair $(\mathcal{H}, \xi)$ carries weight $(-2, +1)$, so $\mathcal{H} \xi dt$ is invariant.

For a finite next touch at $T > t$, define

$$q_\kappa(t, T) \equiv \frac{\kappa}{e^{\kappa(T-t)} - 1}, \quad \lim_{\kappa \rightarrow 0} q_\kappa(t, T) = \frac{1}{T-t}. \quad (4.46)$$

The transformed first-touch density has logarithmic state derivative

$$\partial_\xi \log h_T(t, \xi) = \frac{1-\theta}{\xi} - q_\kappa(t, T). \quad (4.47)$$

Therefore the conditional joint process has drifts

$$\mathbb{E}[d\xi_t | T_{\text{touch}} = T, \mathcal{F}_t] = (2 - \theta - \kappa \xi_t - 2q_\kappa(t, T)\xi_t) dt, \quad (4.48)$$

$$\frac{\mathbb{E}[dS_t | T_{\text{touch}} = T, \mathcal{F}_t]}{S_t} = \rho \sqrt{\mathcal{H}} ((1-\theta) - q_\kappa(t, T)\xi_t) dt. \quad (4.49)$$

These are state functions under the touch-conditioned measure, not averages over ξ . They reduce to the zero-mean-reversion formulae above when $\kappa \rightarrow 0$. The terminal conditional law of S_T and the joint bridge given both T and S_T are collected in Appendix B.

In particular the never-touch Girsanov drift takes the mean-reversion-covariant form

$$\mathbb{E}[d \log S \mid \forall_t \xi_t > 0] - \mathbb{E}[d \log S] = \rho \sqrt{\tilde{\mathcal{H}}} (1-\theta) d\tilde{t} = \rho \sqrt{\mathcal{H}} (1-\theta) e^{-\kappa t} d\tilde{t} = \rho \sqrt{\mathcal{H}} (1-\theta) dt, \quad (4.50)$$

i.e. the drift on $\log S$ is invariant under Eq. 4.41, as it must be for a scalar. Finite- t_∞ conditioning is obtained by composing Eq. 4.41 with the scale-invariant inversion Eq. 4.33, yielding the mean-reverting analogue of Eq. 4.38.

In calculating the Laplace transform for the distribution of the time elapsed before returning to the boundary, first note that

$$e^{\lambda(t-t_0)} = \left(\frac{\kappa \xi}{Z(t_0, \xi; t)} + 1 \right)^{\frac{\lambda}{\kappa}}.$$

Therefore,

$$\begin{aligned}\psi_\xi(\lambda) &= - \int_0^\infty e^{\lambda \frac{\xi}{z}} d\Omega(z) = \frac{1}{\Gamma(1-\theta)} \int_0^\infty z^{-\theta} e^{-z} \left(\frac{\kappa \xi_0}{z} + 1 \right)^{\frac{\lambda}{\kappa}} dz \\ &= \frac{\Gamma(-\theta - \frac{\lambda}{\kappa} + 1) U(-\frac{\lambda}{\kappa}, \theta, \kappa \xi)}{\Gamma(1-\theta)}.\end{aligned}\quad (4.51)$$

Here, U denotes the Tricomi confluent hypergeometric function. Performing an expansion of this expression in the limit $\xi \rightarrow 0$ obtain

$$\psi_\xi(\lambda)_{\xi \rightarrow 0} = 1 + \frac{\Gamma(\theta-1)\Gamma(-\theta - \frac{\lambda}{\kappa} + 1)}{\Gamma(1-\theta)\Gamma(-\frac{\lambda}{\kappa})} (\kappa \xi)^{1-\theta} + \mathcal{O}(\xi) \quad (4.52)$$

Substituting this into the formula Eq. 3.7, we can obtain the Laplace exponent generator for the physical time process in the edge time

$$\Psi(\lambda) = \lim_{\xi \rightarrow 0} \xi^{\theta-1} \log \psi_\xi(\lambda) = \frac{\Gamma(\theta-1)\Gamma(-\theta - \frac{\lambda}{\kappa} + 1)}{\Gamma(1-\theta)\Gamma(-\frac{\lambda}{\kappa})} \kappa^{1-\theta}. \quad (4.53)$$

Note that unlike in the mean reversion case, the generator $\Psi(\lambda)$ is regular at $\lambda = 0$, and, therefore, all the cumulants of the Levy process are finite.

Local time and boundary conditions

Since mean reversion does not contribute to drift near the boundary, it should also not influence the class of equivalence of the measure associated with the process and, therefore, the definition of edge time. As expected, we can select the stationary martingale solution with the same asymptotic behavior as in the mean reversion-free case. That is,

$$\mathfrak{L} = \frac{1}{1-\theta} \int_0^\xi e^{\kappa \xi} \xi^{-\theta} d\xi = \xi^{\theta-1} + \frac{\theta}{2!(1-\theta)} \kappa \xi^\theta + \frac{\theta+1}{3!(1-\theta)} \kappa^2 \xi^{\theta+1} + \dots, \quad (4.54)$$

so that

$$\mathfrak{L}|_{\xi \rightarrow 0} = \xi^{\theta-1} + o(\xi). \quad (4.55)$$

Consequently, the asymptotic decompositions Eq. 3.15, Eq. 3.22 remain intact. Neither changes in the first order the normalization factor in Eq. 4.5

$$\varepsilon = \frac{\theta(1-\theta)}{\xi'^\theta} (1 + \mathcal{O}(\xi)). \quad (4.56)$$

This confirms that the edge time, indeed, does not depend on the mean reversion level.

Expected increase in edge time

Let us use the formula Eq. 4.19 for the model with mean reversion:

$$\begin{aligned}
\mathbb{E}[\delta \mathfrak{r} | \xi_{t_0} = \xi_0] &= \int_{t_0}^{t'} {}_1G(t|t_0, \xi_0) dt \\
&= \int_{t_0}^{t'} \frac{e^{-Z(\xi_0, t_0; t)}}{-\Gamma(\theta - 1) Y(t, t_0)^\theta} dt. \\
&= -\frac{\xi_0^{1-\theta}}{\Gamma(\theta - 1)} \int_{\frac{\kappa \xi_0}{e^{\kappa(t'-t)} - 1}}^{\infty} \frac{e^{-z}}{(z + \kappa \xi_0)^{1-\theta}} \frac{dz}{z}. \\
&= \xi_0^{1-\theta} \frac{\Gamma(\theta - 1, \frac{\kappa \xi_0}{e^{\kappa(t'-t)} - 1})}{-\Gamma(\theta - 1)} (1 + \mathcal{O}(\kappa \xi_0)).
\end{aligned}$$

Here, we applied Eq. 4.25 together with Eq. 4.44, and Eq. 4.43, and changed the variable under the integral to

$$z = Z(t_0, \xi_0; t) = \frac{\kappa \xi_0}{e^{\kappa(t-t_0)} - 1}.$$

Although integration cannot be performed in closed form, it can be approximated as a power expansion over the mean reversion speed κ .

4.1 Simulation

We propose to generate the process variable ξ along with the edge time $\mathfrak{r}$ by exploiting the joint-distribution decomposition of Appendix A. The scheme is exact for ξ and for the next-contact time; only the edge-time increment on a completed excursion is approximated by its conditional expectation. Mean reversion is handled exactly through the change of variables Eq. 4.41. Let $\{t_i\}_{i=0}^N$ be a time grid selected for simulation, with t_0 pointing to the current time. We specify the simulation state with the following variables:

- $\xi_i \equiv \xi_{t_i}$ - the value of the simulated process
- $\mathfrak{r}_i \equiv \mathfrak{r}_{t_i}$ - accrued edge time
- $T_i \equiv \inf_{t > t_i \& \xi_t = 0} t$ - projected time of the next contact with the boundary

The simulation progresses as follows:

- Start by setting $\mathfrak{r}_0 = 0$ and generating T_0 . According to Eq. 4.9, this can be done by generating $z = Z(t_0, \xi_0, T_0)$ via the gamma distribution of the order $1 - \theta$ and then inverting the formulae Eq. 4.43, Eq. 4.44:

$$T_0 = t_0 + \frac{1}{\kappa} \log \left(\frac{\kappa \xi_0}{z} + 1 \right).$$

This is the same as the central χ^2 distribution of the dimension $2(1 - \theta)$ for the variable $2z$.

- If $T_{i-1} > t_i$, the excursion continues past t_i . The projected boundary contact time and the edge time both carry over:

$$T_i = T_{i-1}, \quad \mathfrak{r}_i = \mathfrak{r}_{i-1}.$$

The next value ξ_i is generated from the conditional process. According to Eq. 4.33, this is the dual affine dynamics with the modified Feller parameter Eq. 4.30 and no mean reversion, expressed in terms of the modified variables:

$$\tilde{\xi}_i = \frac{\xi_i}{|Y(T_i, t_i)|^2} \left(\frac{d\tilde{t}}{dt} \right)^{-1}, \quad \tilde{t}_i = \frac{1}{|Y(T_i, t_i)|}.$$

For $\kappa = 0$ this reduces to $\tilde{\xi}_i = \xi_i / (T_i - t_i)^2$ and $\tilde{t}_i = 1 / (T_i - t_i)$; the modified time interval is

$$\delta\tilde{t} \equiv \tilde{t}_i - \tilde{t}_{i-1} = \frac{1}{T_{i-1} - t_i} - \frac{1}{T_{i-1} - t_{i-1}}.$$

The variable $2\tilde{\xi}_i / \delta\tilde{t}$ then follows a non-central χ^2 distribution of dimension $2\tilde{\theta}$ and non-centrality parameter $2\tilde{\xi}_{i-1} / \delta\tilde{t}$, which can be sampled exactly.

- Otherwise, if $T_{i-1} < t_i$, the excursion starting at T_{i-1} has completed before t_i . We first draw the last hit of the boundary in the interval (T_{i-1}, t_i) from the Dynkin–Lamperti arcsine law Eq. A.2, working in the mean-reversion-adjusted time of Eq. 4.41:

$$V \equiv \frac{\tilde{u} - \tilde{T}_{i-1}}{\tilde{t}_i - \tilde{T}_{i-1}} \sim \text{Beta}(1 - \theta, \theta),$$

from which u is recovered via $\tilde{t} = e^{\kappa t} / \kappa$. Given u , the value ξ_i is drawn from the meander density Eq. A.4 at excursion age $t_i - u$, i.e., the exponential distribution of rate $1/(t_i - u)$ (or, equivalently for $\kappa \neq 0$, a gamma variable of shape one multiplied by the factor $Y(t_i, u) = 1 - e^{\kappa(u - t_i)}$). The edge

time increment over (T_{i-1}, u) is approximated by its conditional expectation Eq. A.6:

$$\mathfrak{r}_i = \mathfrak{r}_{i-1} + \frac{\Gamma(1-\theta)\Gamma(2-\theta)}{\Gamma(\theta)\Gamma(2-2\theta)} \left(\tilde{u} - \tilde{T}_{i-1} \right)^{1-\theta}.$$

Finally, we generate the next contact in the same way as at the beginning of the simulation:

$$T_i = t_i + \frac{1}{\kappa} \log \left(\frac{\kappa \xi_i}{z} + 1 \right),$$

with z drawn from the gamma distribution of order $1 - \theta$.

Applications and Extensions

SDE in a covariant form

Instead of expressing SDEs in terms of standard Brownian processes, we prefer to formulate them in terms of expected drifts and increments of quadratic variation and covariation. While less explicit, this approach is independent of the choice of random factors. For instance, in this format, the standard Heston model can be formulated as follows

$$\begin{aligned} \mathbb{E}[d\xi] &= (\theta - \kappa\xi)dt, & \mathbb{E}[dS] &= 0, \\ d\langle \xi \rangle &= 2\xi dt, & d\langle S \rangle &= \frac{1}{2}S^2\mathcal{H}\xi dt, & d\langle \xi, S \rangle &= \sqrt{\mathcal{H}}S\rho\xi dt. \end{aligned} \quad (4.57)$$

This model can be trivially extended by a constant shift of the underlying variance process

$$\begin{aligned} \mathbb{E}[d\xi] &= (\theta - \kappa\xi)dt, & \mathbb{E}[dS] &= 0, \\ d\langle \xi \rangle &= 2\xi dt, & d\langle S \rangle &= \left(\frac{\mathcal{H}}{2}\xi + B \right) S^2 dt, & d\langle \xi, S \rangle &= \sqrt{\mathcal{H}}S\rho\xi dt. \end{aligned} \quad (4.58)$$

This is the standard way to add calibration flexibility to the Heston model while keeping it affine.

Heston Model with a Quasi-Rough Extension

For the purpose of this study, we ignore rates and dividends and model the asset value S as a martingale process. The results can be trivially extended to include

deterministic rates and proportional dividends². Consider the following pseudo-rough extension of the process Eq. 4.58

$$\begin{aligned}
\mathbb{E}[d\xi] &= (\theta - \kappa\xi)dt, & \mathbb{E}[dS] &= 0, \\
\mathbb{E}[dB] &= (\alpha_B - \beta_B B)dt + (\aleph_B - \beth_B B)d\mathfrak{r}, \\
d\langle S \rangle &= S^2 \left(\left(\frac{\mathcal{H}}{2}\xi + B \right) dt + \mathfrak{v}_S d\mathfrak{r} \right), \\
d\langle \xi \rangle &= 2\xi dt, & d\langle \xi, S \rangle &= \sqrt{\mathcal{H}}S\rho\xi dt, \\
d\langle B \rangle &= B^2\mathfrak{v}_B d\mathfrak{r}, & d\langle B, S \rangle &= BS\mathfrak{v}\sqrt{\mathfrak{v}_B\mathfrak{v}_S}d\mathfrak{r}.
\end{aligned} \tag{4.59}$$

The drift of B has two absolute components ($\alpha_B dt$ in physical time, $\aleph_B d\mathfrak{r}$ in edge time) and two relative components ($-\beta_B B dt$ and $-\beth_B B d\mathfrak{r}$); $\mathfrak{v}_S$, $\mathfrak{v}_B$, and $\mathfrak{v}$ are the rough contributions to $\langle S \rangle$, $\langle B \rangle$, and the correlation between them respectively. The absolute coefficients must be non-negative, $\alpha_B \geq 0$ and $\aleph_B \geq 0$, to prevent B from becoming negative. The relative coefficients $\beta_B, \beth_B$ are unrestricted in sign individually, but their combined mean-reversion contribution $\beta_B dt + \beth_B d\mathfrak{r}$ should be positive on average to preserve process stability. For fixed edge-time increments, the smooth physical-time drift is propagated exactly rather than with Euler discretization. Starting from B_a at t_a , it gives

$$B(t_b) = e^{-\beta_B(t_b-t_a)}B_a + \frac{\alpha_B}{\beta_B} \left(1 - e^{-\beta_B(t_b-t_a)} \right), \quad (\beta_B > 0), \tag{4.60}$$

with the continuous limit $B(t_b) = B_a + \alpha_B(t_b - t_a)$ when $\beta_B = 0$. In the extreme infinite edge-reversion model, the edge-time dynamics pin B to its edge-clock equilibrium $\mathfrak{y}_B = \aleph_B/\beth_B$ after the last touch t_{out} . The value at the next fixing is therefore obtained exactly from Eq. 4.60 with $B_a = \mathfrak{y}_B$ and $t_a = t_{\text{out}}$:

$$B(t_{\text{next}}) = e^{-\beta_B(t_{\text{next}}-t_{\text{out}})}\mathfrak{y}_B + \frac{\alpha_B}{\beta_B} \left(1 - e^{-\beta_B(t_{\text{next}}-t_{\text{out}})} \right). \tag{4.61}$$

For the rough contribution to exist, the boundary must be attainable, which constrains the choice of θ to the interval between 0 and 1. In this settings fair variances can be calculated in closed form, which is essential for handling VIX options. The approach can be further extended by allowing for more general relations between the ξ process and equity volatility. However, this would complicate the handling of VIX products.

Simulation To describe the simulation, we formulate the model in traditional SDE form. For this purpose we need two types of Wiener drivers: W_t^ξ, W_t^S in the

²Proportionality does not preclude discreteness.

physical time, and $W_{\mathfrak{t}}^B, W_{\mathfrak{t}}^S$ in the edge time, such that

$$d\langle W_t^\xi \rangle = d\langle W_t^S \rangle = dt, \quad d\langle W_{\mathfrak{t}}^B \rangle = d\langle W_{\mathfrak{t}}^S \rangle = d\mathfrak{t}.$$

The method of simulating the edge time $\mathfrak{t}$ is described in subsection 4.1. The processes in Eq. 4.59 can be expressed in terms of those Wiener drivers as follows:

$$\begin{aligned} d\xi &= (\theta - \kappa\xi) dt + \sqrt{2\xi} dW_t^\xi, \\ dB &= (\alpha_B - \beta_B B) dt + (\mathfrak{N}_B - \mathfrak{J}_B B) d\mathfrak{t} + B\sqrt{\mathfrak{P}_B} dW_{\mathfrak{t}}^B, \\ dS &= S \left(\sqrt{\frac{\mathcal{H}}{2}} \xi + B dW_t^S + \sqrt{\mathfrak{P}_S} dW_{\mathfrak{t}}^S \right). \end{aligned} \quad (4.62)$$

Then the correlation between the Wiener drivers must be

$$d\langle W_t^\xi, W_t^S \rangle = \frac{\rho}{\sqrt{1 + \frac{2B}{\mathcal{H}\xi}}} dt, \quad d\langle W_{\mathfrak{t}}^B, W_{\mathfrak{t}}^S \rangle = \mathfrak{N} d\mathfrak{t}. \quad (4.63)$$

Discretization ordering Because S is correlated with ξ through the ρ channel, the touch-time / meander sampler for ξ cannot be paired with an independent unconditional Euler step for S . The current scheme generates ξ_{i+1} first and then uses the sampled $\Delta\xi$ in the correlated term of the physical-time log- S increment. The $\mathcal{H}$ channel is applied on every step, including a touch-completing step. This is the required local conditioning structure, but it does not by itself establish that the full joint transition is martingale preserving. A stable numerical martingale bias requires the endpoint-normalized joint bridge described in Appendix B; the present stepwise asset transition should therefore be treated as an approximation pending that construction.

Hybrid Pseudo-Rough Reciprocal Affine Model

We combine the calendar-time Heston contribution with a reciprocal affine volatility state ζ whose diffusion runs only in the rough clock. The hybrid model retains reciprocal contributions to asset variance in both calendar time and rough time. This separates roughness of volatility from roughness of the asset itself and permits their strengths to be tested jointly. Setting either asset weight to zero gives a nested specialization: in particular, the calendar-only asset has absolutely continuous quadratic variation while its variance still depends on a rough-clock state. The earlier rough-asset reciprocal specification is also retained as a special case.

The state ξ generates the normalized edge clock $\mathfrak{t}_t$ as in subsection 4.1. This clock is continuous, nondecreasing, finite on finite calendar horizons, and increases

only when $\xi = 0$. The parameters θ and ϑ are the normalized Feller characteristics of ξ and ζ , respectively, under domestic cash. The Heston coefficient $\mathcal{H}$ and correlation ρ retain their meaning from (4.59); $\mathcal{H}$ is distinct from a Hurst exponent.

All coefficients below are constant. The Greek coefficients α and β govern calendar drift of ζ ; $\mathfrak{I}$ governs its rough linear drift, and $\mathfrak{V}_\zeta$ is its rough diffusion variance coefficient. Nonnegative weights c_c and $\mathfrak{V}_S$ specify the calendar and rough reciprocal contributions to asset variance. The subscripts on $\mathfrak{V}_S$ and $\mathfrak{V}_\zeta$ distinguish the stock and reciprocal-state rough variance coefficients. The correlation $\mathfrak{N}$ couples the rough shocks of ζ and S and is relevant only when $\mathfrak{V}_S > 0$. We use a discounted asset, with deterministic rates and dividends suppressed, and write $\mathbb{E}_d[dX]$ for its local drift characteristic under the domestic cash measure $\mathbb{Q}^d$.

$$\begin{aligned}
\mathbb{E}_d[d\xi] &= (\theta - \kappa\xi) dt, \\
\mathbb{E}_d[d\zeta] &= (\alpha - \beta\zeta) dt + \left(\frac{1}{2}\mathfrak{V}_\zeta\vartheta - \mathfrak{I}\zeta\right) d\mathfrak{r}, \\
\mathbb{E}_d[dS] &= 0, \\
d\langle\xi\rangle &= 2\xi dt, \\
d\langle\zeta\rangle &= \mathfrak{V}_\zeta\zeta d\mathfrak{r}, \\
d\langle S\rangle &= S^2 \left[\left(\frac{\mathcal{H}}{2}\xi + \frac{c_c}{\zeta}\right) dt + \frac{\mathfrak{V}_S}{\zeta} d\mathfrak{r} \right], \\
d\langle\xi, S\rangle &= \rho S\sqrt{\mathcal{H}}\xi dt, \\
d\langle\zeta, S\rangle &= \mathfrak{N}S\sqrt{\mathfrak{V}_\zeta\mathfrak{V}_S} d\mathfrak{r}, \\
d\langle\xi, \zeta\rangle &= 0.
\end{aligned} \tag{4.64}$$

The calendar reciprocal stock shock is independent of the Heston and reciprocal-state drivers. Calendar and edge-time martingale brackets are mutually singular, so this shock introduces no additional instantaneous covariation with ζ . The two rough diffusion amplitudes are $S\sqrt{\mathfrak{V}_S/\zeta}$ and $\sqrt{\mathfrak{V}_\zeta\zeta}$, giving the last nonzero covariation above. The correlation of $1/\zeta$ with the rough stock shock has sign opposite to $\mathfrak{N}$.

Calendar-time and rough-time asset specializations. The calendar-time stock model has $c_c = 1$ and $\mathfrak{V}_S = 0$:

$$\frac{d\langle S\rangle}{S^2} = \left(\frac{\mathcal{H}}{2}\xi + \frac{1}{\zeta}\right) dt, \quad d\langle\zeta, S\rangle = 0. \tag{4.65}$$

Here volatility has a rough-clock component, but asset quadratic variation is absolutely continuous in calendar time. Reciprocal stock variance continues to accumu-

late during excursions on which the edge clock is flat. The parameter $\mathfrak{D}$ disappears, and no reciprocal conditional stock-return tilt is required. There remains dependence through ξ and the clock it generates; zero instantaneous covariation does not imply independence of future returns and volatility.

The original rough-asset model is the specialization $c_c = 0$, $\mathfrak{V}_S = 1$ in its chosen numerical time and state units. Its simulation is retained in Appendix F; the separate calendar-time stock implementation is described in Appendix G. The hybrid specification allows both weights to be positive.

Units must be fixed before transferring parameter values between these specializations. If calendar time has dimension T and $[\mathfrak{r}] = T^{1-\theta}$, the normalization $c_c = 1$ gives

$$\begin{aligned} [\xi] = [\zeta] = T, \quad [\alpha] = 1, \quad [\kappa] = [\beta] = T^{-1}, \quad [\mathcal{H}] = T^{-2}, \\ [c_c] = 1, \quad [\mathfrak{V}_S] = [\mathfrak{V}_\zeta] = T^\theta, \quad [\mathfrak{z}] = T^{-(1-\theta)}. \end{aligned} \quad (4.66)$$

Thus $\mathfrak{V}_S = 1$ in this normalization denotes a unit-bearing reference coefficient. In the original rough-asset normalization, where that coefficient is dimensionless, $[\zeta] = T^{1-\theta}$ instead. Raw fitted reciprocal parameters cannot be transferred between the two models without the corresponding change of units and variance normalization. When $c_c > 0$, rescaling ζ (including ζ_0), α , $\mathfrak{V}_\zeta$, and $\mathfrak{V}_S$ by $1/c_c$ sets $c_c = 1$, avoiding a redundant fitted scale. The rough absolute drift $\mathfrak{N} = \mathfrak{V}_\zeta \vartheta/2$, when specified independently, must be rescaled by the same factor. With $c_c = 1$, the normalized hybrid stock variance is

$$\frac{d\langle S \rangle}{S^2} = \left(\frac{\mathcal{H}}{2} \xi + \frac{1}{\zeta} \right) dt + \frac{\mathfrak{V}_S}{\zeta} d\mathfrak{r}.$$

Calibration must nevertheless retain the pure-rough stock benchmark. After fixing common numerical time and state units, the transformation

$$(\zeta_0, \alpha, \mathfrak{N}, \mathfrak{V}_\zeta, \mathfrak{V}_S) \mapsto \lambda(\zeta_0, \alpha, \mathfrak{N}, \mathfrak{V}_\zeta, \mathfrak{V}_S), \quad \lambda > 0, \quad (4.67)$$

with β , $\mathfrak{z}$, and all clock and Heston parameters unchanged, scales the entire reciprocal state by λ while preserving the rough stock channel and its correlation. At fixed $c_c = 1$, its calendar reciprocal variance contribution is instead divided by λ . Thus the stock law of the pure-rough reciprocal model is approached as $\lambda \rightarrow \infty$ in the positive-state regime; it is a boundary limit of the normalized family, not an exact finite-parameter submodel with $c_c = 1$. Keeping $c_c = 0$ explicitly available provides the exact benchmark, and finite parameter bounds must not inadvertently rule out its approximation. Fit comparisons require the same quotes, objective, and numerical accuracy; this extension neither guarantees a strict improvement nor ensures that a local optimizer finds the boundary regime. Low- θ profiles must also

remain admissible: for example, $\theta = 0.05$ gives nominal $H = (1 - \theta)/2 = 0.475$, but this near-Brownian scaling alone does not make the clock deterministic or establish a global calendar-time $3/2$ limit.

Cash and asset numeraires. Write $\mathbb{Q}^f$ for the foreign/asset-numeraire measure, with S interpreted as the value of the foreign numeraire in domestic cash. Localizing before any singular boundary, the density S/S_0 gives

$$\mathbb{E}_f[dX] = \mathbb{E}_d[dX] + \frac{d\langle X, S \rangle}{S}. \quad (4.68)$$

Consequently, for $\boldsymbol{\Psi}_\zeta > 0$,

$$\begin{aligned} \mathbb{E}_f[d\zeta] &= (\alpha - \beta\zeta) dt + \left(\tfrac{1}{2}\boldsymbol{\Psi}_\zeta\vartheta + \boldsymbol{\mathfrak{N}}\sqrt{\boldsymbol{\Psi}_\zeta\boldsymbol{\Psi}_S} - \boldsymbol{\mathfrak{I}}\zeta\right) d\mathfrak{r} \\ &= (\alpha - \beta\zeta) dt + \left(\tfrac{1}{2}\boldsymbol{\Psi}_\zeta\vartheta_f - \boldsymbol{\mathfrak{I}}\zeta\right) d\mathfrak{r}, \end{aligned} \quad (4.69)$$

where

$$\boxed{\vartheta_d = \vartheta, \quad \vartheta_f = \vartheta + 2\boldsymbol{\mathfrak{N}}\sqrt{\frac{\boldsymbol{\Psi}_S}{\boldsymbol{\Psi}_\zeta}}.} \quad (4.70)$$

The plus sign reflects the convention that $\boldsymbol{\mathfrak{N}}$ correlates the inverse-variance state ζ with the rough stock shock. The calendar weight c_c produces no direct Feller shift. In particular, the two local reciprocal Feller characteristics coincide when $\boldsymbol{\Psi}_S = 0$. The clock state nevertheless has foreign drift

$$\mathbb{E}_f[d\xi] = \left[\theta - (\kappa - \rho\sqrt{\mathcal{H}})\xi\right] dt. \quad (4.71)$$

Thus the measure change leaves θ unchanged but changes the law of future clock increments. Even when the reciprocal Feller characteristics coincide, the marginal laws of ζ need not agree under the two measures.

Admissibility under both numeraires. For the positive rough-diffusion case we select the interior regime

$$\begin{aligned} 0 < \theta < 1, \quad \kappa \geq 0, \quad \vartheta_d > 1, \quad \vartheta_f > 1, \\ \alpha \geq 0, \quad \boldsymbol{\Psi}_\zeta > 0, \quad \mathcal{H} \geq 0, \quad c_c, \boldsymbol{\Psi}_S \geq 0, \quad |\rho| \leq 1, \quad |\boldsymbol{\mathfrak{N}}| \leq 1, \\ \xi_0 \geq 0, \quad \zeta_0 > 0, \quad S_0 > 0. \end{aligned} \quad (4.72)$$

Equivalently, the reciprocal restrictions are

$$\vartheta > \max \left\{ 1, 1 - 2\boldsymbol{\mathfrak{N}}\sqrt{\frac{\boldsymbol{\Psi}_S}{\boldsymbol{\Psi}_\zeta}} \right\}. \quad (4.73)$$

The clock state can reach zero while the reciprocal state avoids zero under both local numeraire dynamics. The nonnegative calendar intercept prevents calendar drift from driving ζ through zero while the clock is flat. Finite signed β and $\mathfrak{I}$ are allowed by positivity; positive values give mean reversion in their respective clocks. The rough CIR nonattainment threshold itself is a Feller characteristic at least one; the strict bounds select an interior regime with finite first inverse moments for the rough affine transitions. These local boundary restrictions are not, by themselves, a proof of the global martingale property of the full hybrid asset. The distinction is discussed for the ordinary 3/2 model in [14].

Locally to this section, write $\aleph$ for the rough absolute drift:

$$\aleph = \frac{1}{2}\psi_\zeta\vartheta, \quad \vartheta_d = \frac{2\aleph}{\psi_\zeta}, \quad \vartheta_f = \frac{2(\aleph + \mathfrak{I}\sqrt{\psi_\zeta\psi_S})}{\psi_\zeta}. \quad (4.74)$$

For $\psi_\zeta > 0$, one can specify either ϑ or $\aleph$, with the other derived. The zero-diffusion, positive-rough-drift limit below instead keeps $\aleph$ fixed while $\psi_\zeta \downarrow 0$ and $\vartheta = 2\aleph/\psi_\zeta \uparrow \infty$; no finite reciprocal Feller parameter describes that limit.

The boundary restrictions can also be imposed directly on drift and diffusion, without dividing by ψ_ζ :

$$2\aleph \geq \psi_\zeta, \quad 2(\aleph + \mathfrak{I}\sqrt{\psi_\zeta\psi_S}) \geq \psi_\zeta. \quad (4.75)$$

For positive diffusion, strict inequalities select the finite-first-inverse-moment regime. In the calendar-stock specialization $\psi_S = 0$ they reduce to the single condition $2\aleph > \psi_\zeta$. When $\psi_\zeta = 0$, any $\aleph \geq 0$ is allowed with a positive initial state and nonnegative calendar intercept; the Feller ratio is unnecessary. Fitting $\aleph$ and ψ_ζ independently therefore includes a positive rough drift with zero diffusion, whereas fixing finite ϑ ties both strengths together. Both $\aleph$ and ψ_ζ have dimension T^θ under the calendar-stock normalization.

Boundary events and currency devaluation. If a reciprocal characteristic is below one, zero is attainable for the corresponding rough affine diffusion. In the two-clock system one must also account for calendar drift and the rough time available on the chosen horizon. For an asset boundary the relevant accumulated variance is

$$\int_0^t \left(\frac{\mathcal{H}}{2}\xi_s + \frac{c_c}{\zeta_s} \right) ds + \psi_S \int_0^t \frac{d\aleph_s}{\zeta_s}. \quad (4.76)$$

When $\psi_S > 0$, divergence of the inverse-state rough integral at the rough affine zero boundary gives the original devaluation mechanism: under domestic cash the stochastic exponential tends to zero, whereas the corresponding event for $1/S$ under the foreign numeraire describes devaluation of domestic cash. The latter is

not explosion of a nonnegative domestic local martingale under its own measure. Foreign boundary events can instead signal loss of mass and a strict domestic local martingale; see [15]. When $\boldsymbol{\Psi}_S = 0$, divergence of an unused rough integral does not settle the asset boundary behavior. The calendar inverse-state integral and any continuation through $\zeta = 0$ must then be examined separately. The interior restrictions above avoid this issue in the model used here.

Leverage and clock activity. Consider the calendar-time stock specialization (4.65), with $\rho < 0$, $\beta > 0$, α/β small, and little or no rough mean reversion $\mathfrak{A}$. A negative stock innovation through the Heston leverage channel tends to raise ξ , delaying the next boundary touch. During the ensuing flat clock interval, the calendar drift draws ζ towards a small level, so its reciprocal variance increases while its rough-clock activity is suppressed. The Heston variance contribution also increases with ξ . A positive correlated innovation has the opposite tendency: ξ falls, future touches become more likely, and the positive rough intercept pushes ζ upwards. This can lower reciprocal variance while making rough fluctuations more prominent. These are conditional leverage tendencies, not pathwise rules for every stock return; independent stock shocks and the current reciprocal state also matter.

The inverse-state drift makes the local drift conditions explicit:

$$\mathbb{E}_d[d(1/\zeta)] = \left(\frac{\beta}{\zeta} - \frac{\alpha}{\zeta^2} \right) dt + \left[\frac{\mathfrak{A}}{\zeta} + \frac{\boldsymbol{\Psi}_\zeta(1 - \vartheta/2)}{\zeta^2} \right] d\mathfrak{R}. \quad (4.77)$$

For $\alpha = \mathfrak{A} = 0$, the calendar drift raises reciprocal variance, while the local rough drift lowers it when $\vartheta > 2$. Merely imposing a positive rough drift of ζ is insufficient: the inverse transformation has an Itô correction. At $1 < \vartheta < 2$ that correction reverses the sign of the rough inverse-state drift. Neither sign alone orders total volatility-of-volatility. Indeed, for $V = \mathcal{H}\xi/2 + 1/\zeta$ in this specialization,

$$d\langle V \rangle = \frac{\mathcal{H}^2 \xi}{2} dt + \frac{\boldsymbol{\Psi}_\zeta}{\zeta^3} d\mathfrak{R}. \quad (4.78)$$

An especially simple example retains only proportional calendar drift and constant rough drift, with no reciprocal diffusion:

$$d\zeta = -\beta\zeta dt + \mathfrak{K} d\mathfrak{R}, \quad d(1/\zeta) = \frac{\beta}{\zeta} dt - \frac{\mathfrak{K}}{\zeta^2} d\mathfrak{R}, \quad \beta, \mathfrak{K} > 0. \quad (4.79)$$

The explicit solution is

$$\zeta_t = e^{-\beta(t-t_0)} \zeta_{t_0} + \mathfrak{K} \int_{t_0}^t e^{-\beta(t-s)} d\mathfrak{R}_s > 0. \quad (4.80)$$

It produces smooth growth of reciprocal variance between touches and clock-driven reductions when touches resume. This limit remains random and can have singular, locally irregular changes, but ζ and $1/\zeta$ are continuous finite-variation processes with zero quadratic variation. Their randomness comes entirely from the clock. The $(1 - \theta)/2$ nominal scaling of a Brownian motion evaluated at rough time therefore does not carry over unchanged: a pure clock drift instead has the clock's own $1 - \theta$ boundary scaling. In either construction, a change in clock activity at fixed θ is not a change of the structural clock exponent; it changes the prominence of irregular behavior over an observation window.

Empirical motivation and its limits. There is empirical support for the qualitative association of lower volatility with greater roughness, and market distress with less roughness. Gatheral, Jaisson, and Rosenbaum [2, Section 2.6] split their realized-variance sample into two halves and observe higher estimated Hurst exponents in the half containing the financial crisis, although estimates remain in the rough range 0.06–0.20. Bennedsen, Lunde, and Pakkanen [3, Section 5.1.1 and Figure 3] study E-mini S&P 500 data from 2005–2014 using noise-robust rolling estimates of volatility roughness. They find periods of greater smoothness during the 2007–2008 financial crisis and the 2011 Greek debt crisis. Their estimates use approximately 150-day windows and support a broad regime association, with the authors cautioning against definitive conclusions from a single time series.

Evidence from option prices points in the same direction. Rømer [4, Section 4.1] calibrates to SPX options over 2004–2019 and reports positive correlations between the fitted Hurst exponent and VIX: 0.51 for rough Bergomi and 0.10 for rough Heston. Higher Hurst exponents mean smoother volatility. The calibration result is model dependent, however: the fractional kernel makes its Hurst parameter also absorb aspects of mean reversion at longer horizons.

These findings motivate the regime dependence of the hybrid model; they do not establish that a positive return directly causes greater roughness, nor that roughness disappears during drawdowns. The proposed clock mechanism is one possible explanation whose strength must be checked against regime-conditioned volatility increments and option surfaces. In particular, the drift-only limit is a distinct hypothesis, not a consequence of the empirical studies. The full hybrid model retains a direct rough stock channel with $\psi_S > 0$ alongside calendar stock variance; the restriction $\psi_S = 0$ supplies a nested benchmark for testing its contribution.

Conditional forward variance. For a window of length Δt , define the idealized conditional variance-swap quantity

$$\begin{aligned} \frac{\text{VIX}_t^2}{100^2} = \frac{1}{\Delta t} & \left[\frac{\mathcal{H}}{2} \mathbb{E}_t \int_t^{t+\Delta t} \xi_s ds + c_c \mathbb{E}_t \int_t^{t+\Delta t} \frac{ds}{\zeta_s} \right. \\ & \left. + \boldsymbol{\Psi}_S \mathbb{E}_t \int_t^{t+\Delta t} \frac{d\mathfrak{r}_s}{\zeta_s} \right], \end{aligned} \quad (4.81)$$

with times measured in years and $\Delta t = 30/365$ for the usual window. The conditional Heston area is explicit:

$$\mathbb{E}_t \int_t^{t+\Delta t} \xi_s ds = \xi_t \frac{1 - e^{-\kappa \Delta t}}{\kappa} + \frac{\theta}{\kappa} \left(\Delta t - \frac{1 - e^{-\kappa \Delta t}}{\kappa} \right), \quad (4.82)$$

with limit $\xi_t \Delta t + \theta \Delta t^2 / 2$ at $\kappa = 0$. When $\boldsymbol{\Psi}_S = 0$, freezing ζ over the window adds c_c / ζ_t to the annualized Heston variance estimate. The shift is the inverse state, not $1/\vartheta$. This is an approximation to the calendar inverse-state expectation, not an exact VIX formula; future rough state dynamics and inverse-state convexity remain relevant.

Scale-Covariant Reciprocal Model (experimental)

As a separate modeling option, set $p = 1/(1 - \theta)$ and normalize the discounted asset to $S_0 = 1$. The proposed domestic local characteristics are

$$\begin{aligned} \mathbb{E}_d[d\xi] &= (\theta - \kappa_d \zeta^{-p} \xi) dt, & d\langle \xi \rangle &= 2\xi dt, \\ \mathbb{E}_d[d\zeta] &= (\vartheta_d - \mathfrak{N}^{1-\theta} \zeta) d\mathfrak{r}, & d\langle \zeta \rangle &= 2\zeta d\mathfrak{r}, \\ \mathbb{E}_d[dS] &= 0, & d\langle \xi, \zeta \rangle &= 0, \\ \frac{d\langle S \rangle}{S^2} &= \Sigma \xi \zeta^{-2p} dt + \frac{\boldsymbol{\Psi}}{\zeta} d\mathfrak{r}, & & \\ \frac{d\langle S, \xi \rangle}{S} &= \rho \sqrt{2\Sigma} \xi \zeta^{-p} dt, & \frac{d\langle S, \zeta \rangle}{S} &= \mathfrak{N} \sqrt{2\boldsymbol{\Psi}} d\mathfrak{r}. \end{aligned} \quad (4.83)$$

Both covariance coefficients Σ and $\boldsymbol{\Psi} = \boldsymbol{\Psi}$ are dimensionless; Σ is a variance coefficient, not a volatility. Here $\boldsymbol{\Psi}$ scales the rough stock channel, while the reciprocal-state quadratic variation has coefficient two. The two stock channels have correlations ρ with ξ and $\mathfrak{N} = \mathfrak{N}$ with ζ . Since ζ is an inverse-variance state, positive $\mathfrak{N}$ corresponds to negative stock/volatility-level correlation in the rough channel.

Units and the two numeraire. This version uses

$$[\xi] = T, \quad [\zeta] = [\mathfrak{r}] = T^{1-\theta}, \quad [\mathfrak{N}] = T^{-1}. \quad (4.84)$$

The parameters $\theta, \kappa, \vartheta, \Sigma, \boldsymbol{\Psi}, \rho, \mathfrak{N}$ are dimensionless. Locally to this subsection, $\mathfrak{N}$ is a positive calendar rate, rather than the rough absolute drift used in the preceding model. The localized drift change (4.68) gives

$$\kappa_{\text{f}} = \kappa_{\text{d}} - \rho\sqrt{2\Sigma}, \quad \vartheta_{\text{f}} = \vartheta_{\text{d}} + \mathfrak{N}\sqrt{2\Psi}. \quad (4.85)$$

The other coefficients are unchanged. For the inverse asset under the foreign numeraire, both correlations reverse sign, making this parameter map involutive.

The numerical experiment selects $\xi_0 \geq 0$, $\zeta_0 > 0$, $\mathfrak{N} > 0$, nonnegative variance coefficients, correlations in $[-1, 1]$, positive domestic and foreign reversion coefficients, and

$$0 < \theta < 1, \quad \vartheta_{\text{d}} > 2p + \max\{0, -\mathfrak{N}\sqrt{2\Psi}\}. \quad (4.86)$$

The last condition is stronger than Feller nonattainment. For a CIR state at deterministic edge time, the stationary inverse moment is $(\mathfrak{N}^{1-\theta})^r \Gamma(\vartheta - r) / \Gamma(\vartheta)$ for $r < \vartheta$. Thus the restriction makes the stationary edge-time marginal moment of order $2p$ finite under both parameter sets. It does not establish calendar-time integrability of $\xi\zeta^{-2p}$ or a true-martingale density for the coupled model.

Simulation and observable-state normalization. The reversion coordinate is accumulated along each path:

$$\lambda_t = \kappa_{\text{d}} \zeta_t^{-p}, \quad \tau_t = \int_0^t \lambda_s ds. \quad (4.87)$$

While the rough clock is flat, ζ and λ are constant. For $G \sim \Gamma(1 - \theta, 1)$, the retained next-contact coordinate satisfies $\tau_{\text{hit}} - \tau_t = \log(1 + \lambda_t \xi_t / G)$. The scheme keeps this coordinate across no-contact steps, uses its conditional endpoint law, and advances ζ through an approximate intrinsic boundary increment when contacts occur. That increment uses a conditional bridge mean and Jacobi quadrature; it is not an exact draw of the local-time path. Calendar substeps also bound $\lambda \Delta t$.

Stock endpoint factors use a noncentral-Gamma cumulant

$$K(b) = -a \log(1 - sb) + \frac{\nu b}{1 - sb}, \quad sb < 1. \quad (4.88)$$

The multiplier $\exp(bX - K(b))$ has unit mean under the law defining K , with its slope selected to match a frozen-step foreign first moment. For the calendar factor,

the cumulant parameters must be those of the CIR endpoint conditional on the current observable state:

$$a = \theta, \quad s = \frac{1 - e^{-\lambda \Delta t}}{\lambda}, \quad \nu = \xi_t e^{-\lambda \Delta t}. \quad (4.89)$$

The internally sampled future contact is not part of the financial filtration. Normalizing instead under its conditional bridge law incorrectly removes the asset weighting of contact probabilities, even when the sample asset mean remains close to one. The implemented observable-state normalization fixes this demonstrated error; a first-contact survival test under the foreign drift verifies the change as the grid is refined. The reciprocal factor uses the rough-time CIR endpoint law. Residual stock shocks are normalized lognormal factors with expected integrated variance and Gaussian antithetics.

The implementation remains a finite-step approximation to the joint model. Neither the survival check nor asset-mean tests prove convergence of the full coupled process or its global continuous-time martingale property. Historical fits reported by another working copy are not adopted as verified evidence: their checkpoints and market observations were not supplied, and the normalization repair changes the simulator. The new module and notebook are kept separate, reuse the existing volume-aware calibration objective and interruption/checkpoint hooks, and retain ThetaData as the default market source. Bloomberg is an optional adapter. CPU and the existing ROCm GPU path remain supported; PyTorch exposes both ROCm and NVIDIA devices through its CUDA API. Further implementation details and validation are recorded in `scale_covariant_reciprocal_model.md`.

Exponential rough-volatility model (experimental)

The confirmed model retains the boundary-state factor. Let $p = (1 - \theta)^{-1}$, $0 < \theta < 1$, and use a forward-normalized stock. Its domestic characteristics are

$$\begin{aligned} V_t &:= S_t^{-2} \frac{d\langle S \rangle_t}{dt} = \Sigma \xi_t e^{\lambda \xi_t} \zeta_t^{-2p}, \\ \mathbb{E}_d[d\xi_t \mid \mathcal{F}_t] &= \left(\theta - \kappa_d \xi_t e^{\lambda \xi_t / 2} \zeta_t^{-p} \right) dt, & d\langle \xi \rangle_t &= 2\xi_t dt, \\ \mathbb{E}_d[d\zeta_t \mid \mathcal{F}_t] &= \left(\vartheta - \aleph^{1-\theta} \zeta_t \right) d\mathfrak{r}_t, & d\langle \zeta \rangle_t &= 2\zeta_t d\mathfrak{r}_t, \\ \mathbb{E}_d[dS_t \mid \mathcal{F}_t] &= 0, & d\langle \xi, \zeta \rangle_t &= 0, \\ \frac{d\langle S, \xi \rangle_t}{S_t} &= \rho \sqrt{2\Sigma} \xi_t e^{\lambda \xi_t / 2} \zeta_t^{-p} dt, & d\langle S, \zeta \rangle_t &= 0. \end{aligned} \quad (4.90)$$

Stock diffusion uses calendar time. The reciprocal state evolves on boundary time $\mathfrak{r}$, supported on $\{\xi = 0\}$, and affects future variance. The dimensions are

$$[\xi] = T, \quad [\zeta] = [\mathfrak{r}] = T^{1-\theta}, \quad [\lambda] = [\aleph] = T^{-1}, \quad [\Sigma] = [\kappa] = [\theta] = [\vartheta] = [\rho] = 1.$$

Thus $\lambda\xi$ is dimensionless and Σ is a variance coefficient. The fixed selector `xi_power=1` identifies this default, rather than a physical calibration parameter.

Numeraire closure and boundary behavior. Matching covariance and nonlinear reversion gives the localized map

$$\kappa_f = \kappa_d - \rho\sqrt{2\Sigma}. \quad (4.91)$$

The inverse stock reverses ρ , with other coefficients unchanged. Since $\int \mathbf{1}_{\{\xi_t=0\}} dt = 0$, calendar and boundary time are mutually singular. Stock quadratic variation has zero cross variation with the rough reciprocal martingale; ϑ has no numeraire shift.

Validation requires $\xi_0 \geq 0$, $\zeta_0 > 0$, $\Sigma \geq 0$, $\varkappa > 0$, $\lambda \geq 0$, $|\rho| \leq 1$, nonnegative domestic and foreign reversion, and $\vartheta > 2p$. The last bound makes the stationary rough-time marginal inverse moment of order $2p$ finite, without establishing calendar-time joint variance integrability or continuous stock martingality.

For frozen positive ζ and $c = \kappa\zeta^{-p} > 0$, the stationary density is

$$\pi(x) \propto x^{\theta-1} \exp\left[-\frac{2c}{\lambda} \left(e^{\lambda x/2} - 1\right)\right], \quad \lambda > 0. \quad (4.92)$$

At $\lambda = 0$ it is Gamma with shape θ and rate c . Positive λ strengthens large-state reversion while preserving the drift intercept θ . The accessible boundary is instantaneously reflecting, with local clock exponent $1 - \theta$, rather than absorbing. Negative foreign reversion and $\lambda > 0$ permit finite-time explosion even with frozen ζ . Nonnegative reversion excludes that mechanism without settling the full coupled model.

The state factor masks direct rough variance. Every rough-clock Itô term in V vanishes because $\xi = 0$ on the clock's support. Consequently,

$$d\langle V \rangle_t = 2\Sigma^2 \zeta_t^{-4p} \xi_t e^{2\lambda\xi_t} (1 + \lambda\xi_t)^2 dt. \quad (4.93)$$

The rough state still changes later excursion amplitudes and conditional variance. The exponential amplifies calendar fluctuations at high ξ , without altering the boundary exponent or restoring direct rough instantaneous variance.

For comparison only, `xi_power=0` selects

$$V^{(0)} = \Sigma e^{\lambda\xi} \zeta^{-p}, \quad \mathbb{E}_d[d\xi \mid \mathcal{F}_t] = \left(\theta - \kappa_d \sqrt{\xi} e^{\lambda\xi/2} \zeta^{-p/2}\right) dt.$$

Its covariance is $\rho\sqrt{2\Sigma\xi} e^{\lambda\xi/2} \zeta^{-p/2} dt$, with the same map, and its direct rough variation survives:

$$\frac{d\langle V^{(0)} \rangle_t}{(V_t^{(0)})^2} = 2\lambda^2 \xi_t dt + \frac{2p^2}{\zeta_t} d\mathfrak{r}_t.$$

This comparison requires $\vartheta > p$, changes reversion and the reciprocal exponent, and is not the confirmed default.

Optional calendar mean reversion of the reciprocal state. An additional extension keeps the stock variance and all quadratic variations above, while replacing only the reciprocal drift by

$$\mathbb{E}_d[d\zeta_t \mid \mathcal{F}_t] = (\alpha - \beta\zeta_t)dt + (\vartheta - \aleph^{1-\theta}\zeta_t)d\mathfrak{r}_t, \quad d\langle\zeta\rangle_t = 2\zeta_t d\mathfrak{r}_t. \quad (4.94)$$

Here $\alpha, \beta \geq 0$, $[\alpha] = T^{-\theta}$ and $[\beta] = T^{-1}$. The choice $\alpha = \beta = 0$ recovers the boundary-time-only reciprocal dynamics. For $\beta > 0$, the calendar subflow reverts towards α/β ; from a positive starting state its exact solution is

$$\Phi_h(z) = \begin{cases} e^{-\beta h}z + \alpha(1 - e^{-\beta h})/\beta, & \beta > 0, \\ z + \alpha h, & \beta = 0. \end{cases} \quad (4.95)$$

This map preserves positivity at every finite $h \geq 0$ and needs no floor. The source term can be evaluated with `expm1` to retain accuracy near $\beta h = 0$. The separate implementation surrounds the original fresh-touch and rough CIR step with two calendar ζ half-flows. It still requires grid refinement; exactness of a subflow is not exactness of the coupled transition. Stock factors retain variance from the original observable pre-step states. Its module is `py/rough_exponential_vol_smooth_zeta.py`, with separate adapter `py/rough_exponential_vol_smooth_zeta_calibration.py`.

The new calendar term has finite variation. It therefore adds no covariance with the stock and no numeraire shift to α or β ; the map (4.91) remains unchanged. The rough quadratic-variation term in V remains masked by $\xi = 0$. Calendar movement of ζ does, however, contribute the smooth finite-variation term

$$dV_t|_{\text{calendar } \zeta} = -2p V_t \frac{\alpha - \beta\zeta_t}{\zeta_t} dt.$$

Equation (4.93) is consequently unchanged. The rough-time Gamma law describes the rough CIR subflow, rather than the stationary law of the full calendar-drift extension. Its moment restriction does not establish joint calendar-time integrability. This extension is a separate model choice; calibration and validation of the original model do not validate its additional drift.

Deterministic flows and fresh splitting. Freeze ζ , put $c = \kappa\zeta^{-p}$ and $a = \lambda/2$, and solve $\dot{x} = -cxe^{ax}$. For $a > 0$, the default flow satisfies

$$E_1(aD_h(x)) = E_1(ax) + ch, \quad E_1(z) = \int_z^\infty \frac{e^{-u}}{u} du. \quad (4.96)$$

At $a = 0$, $D_h(x) = xe^{-ch}$. Positive states remain positive at every finite substep; zero remains zero. Stable integrated-clock inverses require no state floor.

The default state scheme applies a drift half-step with current ζ , a fresh full transition for $d\xi = \theta dt + \sqrt{2\xi} dW$, a reciprocal CIR update on the approximate boundary clock, and a drift half-step with the new ζ . Symmetry alone does not prove second-order convergence of the full singular-clock, stock and variance approximation.

From diffusion starting value $\bar{\xi}$, the contact draw is $H = \bar{\xi}/G$, $G \sim \Gamma(1 - \theta, 1)$, used only for that transition. A prior forecast is not retained after nonlinear drift changes the state. The endpoint marginal is exact in ideal arithmetic. From zero it is $h\Gamma(\theta, 1) > 0$ almost surely, and the final default half-flow preserves positivity. The clock increment replaces random bridge local time by a conditional mean; its full law and joint law with the endpoint remain approximate. The reciprocal transition uses that approximate clock.

For the optional literal branch, $c = \kappa\zeta^{-p/2}$ and $\dot{x} = -c\sqrt{x}e^{ax}$. With $F_a(x) = \sqrt{\pi/a} \operatorname{erf}(\sqrt{ax})$ for $a > 0$,

$$F_a(D_h(x)) = \max\{F_a(x) - ch, 0\}.$$

The deterministic extinction time is $F_a(x)/c$; at $a = 0$ the flow is $\max\{\sqrt{x} - ch/2, 0\}^2$. This branch puts the full drift before fresh diffusion, avoiding an artificial zero atom at the final endpoint.

A formal default integrating-factor alternative would use

$$K_t = \int_0^t \kappa e^{\lambda\xi_s/2} \zeta_s^{-p} ds, \quad \tilde{\xi}_t = e^{K_t} \xi_t, \quad d\tilde{t} = e^{K_t} dt.$$

This clock is path-dependent. The literal rate instead contains $\xi_s^{-1/2} e^{\lambda\xi_s/2} \zeta_s^{-p/2}$, singular near zero. Global finiteness and boundary behavior need justification. Neither branch implements an integrating-factor sampler; both use fresh splits.

Predictable stock factors and conditional VIX. The stock multiplier uses observable pre-step states:

$$v_i = V(\xi_i, \zeta_i) h_i, \quad \frac{S_{i+1}}{S_i} = \exp \left[-\frac{1}{2} v_i + \sqrt{v_i} \left(\rho Z_i + \sqrt{1 - \rho^2} Z_i^\perp \right) \right]. \quad (4.97)$$

The shared Gaussian is independent of touch/Beta priors, and the residual Gaussian is independent. This gives conditional mean one in exact arithmetic without normalization under a hidden future contact. Sharing an endpoint Gaussian still only approximates continuous joint dynamics. Discrete martingality does not prove continuous martingality, numeraire duality or clock convergence. Grid refinement

and independent seeds remain necessary; antithetic errors use independent quartet means.

For a calendar window Δ , the confirmed model defines

$$M_{\Delta}(\xi_t, \zeta_t) = \frac{1}{\Delta} \mathbb{E}_d \left[\int_t^{t+\Delta} \Sigma \xi_u e^{\lambda \xi_u} \zeta_u^{-2p} du \middle| \mathcal{F}_t \right], \quad \text{VIX}_t = 100 \sqrt{M_{\Delta}(\xi_t, \zeta_t)}. \quad (4.98)$$

Both future states evolve in this expectation: rough-state effects survive instantaneous masking at contact. Freezing ζ or inserting an affine-Heston expectation defines a different proxy. The helper uses nested simulation with $\Delta = 30/365$ years by default. Its inner mean targets the finite-grid split model; a square root introduces finite-inner-sample Jensen bias: $\mathbb{E}[\sqrt{\widehat{M}} \mid \xi, \zeta] \leq \sqrt{M}$ for an unbiased inner estimate. The option-price bias has no universal sign. Check inner path counts, grids and independent seeds before interpreting VIX prices. Convex variance and faster reversion do not guarantee positive VIX implied skew or a joint SPX/VIX fit.

The implementation is `py/rough_exponential_vol.py`, with adapter `py/rough_exponential_vol_calibration.py` and notebook `py/example_exponential_rough_vol_calibration.ipynb`. Primary references are Baldeaux and Badran’s nonlinear pure-diffusion VIX example, Alòs, García-Lorite and Muguruza’s VIX skew analysis, and Mijatović and Urusov’s diffusion martingale criteria. These motivate the experiment and its checks; they do not establish its coupled mathematical or calibration properties.

Reciprocal affine dynamics in a mixed clock

An alternative separates the generation of operational time from the stock and volatility dynamics. The affine state ξ generates the canonical boundary clock $\mathfrak{r}$, and both the discounted stock and its reciprocal volatility state ζ run in the same mixed clock

$$\tau_t = t + \boldsymbol{\Psi}_{\tau} \mathfrak{r}_t, \quad \boldsymbol{\Psi}_{\tau} \geq 0. \quad (4.99)$$

For calibration, t denotes our normalized business time. A quoted calendar maturity T maps to $t = b(T)$ through a deterministic activity clock. Weekends and full-day NYSE holidays receive a rate approximately 0.103214 times the weekday rate, and a common normalization makes the mean rate over the full year equal to one. The current implementation uses the 2026 NYSE holiday calendar. Longer-expiry experiments extend the published holiday calendar through 2027 while retaining the same 2026 normalization and all earlier clock values. Both the affine clock generator and the reciprocal stock and volatility simulation use this business time.

Quoted option maturity T , carry, discount factors and displayed implied volatility retain their calendar-time conventions.

The coefficient of business time is fixed at one. The Brownian pair driving the stock and reciprocal state in operational time is independent of the affine clock driver. Thus the clock can be generated first, before the reciprocal stock model is simulated on its path-dependent τ grid.

Here ρ denotes the correlation of stock returns and variance, rather than the correlation with the inverse-variance state. In particular, negative ρ represents negative volatility–stock correlation, while the correlation of ζ and stock innovations is $-\rho$. The complete drift and covariance specification is

$$\begin{aligned} \mathbb{E}_d[d\xi] &= (\theta - \kappa\xi) dt, & d\langle\xi\rangle &= 2\xi dt, \\ d\tau &= dt + \boldsymbol{\Psi}_\tau d\mathfrak{t}, & \mathbb{E}_d[dS] &= 0, \\ \mathbb{E}_d[d\zeta] &= (\tfrac{1}{2}\boldsymbol{\Psi}_\zeta\vartheta - \mathfrak{I}\zeta) d\tau, & d\langle\zeta\rangle &= \boldsymbol{\Psi}_\zeta\zeta d\tau, \\ d\langle S\rangle &= \frac{S^2}{\zeta} d\tau, & d\langle\zeta, S\rangle &= -\rho S\sqrt{\boldsymbol{\Psi}_\zeta} d\tau, \\ d\langle\xi, \zeta\rangle &= 0, & d\langle\xi, S\rangle &= 0. \end{aligned} \tag{4.100}$$

In business-time coordinates, the same reciprocal state controls both stock variance contributions: $S^2\zeta^{-1}dt$ and $S^2\boldsymbol{\Psi}_\tau\zeta^{-1}d\mathfrak{t}$. There is no separate Heston stock component or additional stock clock. The parameters $\boldsymbol{\Psi}_\tau$ and $\boldsymbol{\Psi}_\zeta$ distinguish clock mixing from reciprocal-state diffusion.

The clock requires $0 < \theta < 1$, $\kappa \geq 0$ and $\xi_0 \geq 0$. For the reciprocal model take $\zeta_0 > 0$, $\boldsymbol{\Psi}_\zeta > 0$, $\mathfrak{I} \geq 0$ and $|\rho| \leq 1$. Its cash and stock-numeraire Feller characteristics are

$$\vartheta_d = \vartheta, \quad \vartheta_f = \vartheta - \frac{2\rho}{\sqrt{\boldsymbol{\Psi}_\zeta}}. \tag{4.101}$$

Nonattainment of zero requires the corresponding characteristic to be at least one. The implementation uses strict inequalities above one, also ensuring finite first inverse-state moments. The clock law is unchanged under the stock numeraire: conditioning on the independent, finite clock preserves the true-martingale property of the underlying reciprocal stock model. Introducing instantaneous stock– ξ correlation would require a different joint simulation and numeraire adjustment.

For $V = 1/\zeta$, Itô’s formula gives the familiar reciprocal, or $3/2$, variance dynamics in operational time:

$$\mathbb{E}_d[dV] = [\mathfrak{I}V + \boldsymbol{\Psi}_\zeta (1 - \tfrac{1}{2}\vartheta) V^2] d\tau, \quad d\langle V\rangle = \boldsymbol{\Psi}_\zeta V^3 d\tau. \tag{4.102}$$

The usual stabilizing quadratic variance drift requires $\vartheta > 2$ and $\mathfrak{I} > 0$, which is stronger than boundary nonattainment. At $\boldsymbol{\Psi}_\tau = 0$ the model reduces to the ordinary smooth reciprocal model run in business time and is independent of the affine clock parameters.

Dimensions and simulation. The states ξ and ζ , and clocks t and τ , have time dimension one; $\mathfrak{t}$ has time dimension $1 - \theta$. Thus $\kappa, \mathfrak{z} \sim \text{time}^{-1}$, $\boldsymbol{\psi}_\tau \sim \text{time}^\theta$, while $\theta, \vartheta, \rho, \boldsymbol{\psi}_\zeta$ are dimensionless. The quantity $\boldsymbol{\psi}_\tau^{1/\theta}$ is a time scale. Calibration may use $\boldsymbol{\psi}_\tau/t_{\text{ref}}^\theta$ to expose a dimensionless mixing strength at a fixed reference maturity expressed in business time.

First build the calendar grid, retaining quoted maturities, and map its points T_i to $t_i = b(T_i)$, integrating the piecewise business-clock rates exactly. Generate a bank of $(\xi_i, \mathfrak{t}_i)$ on this internal business-time grid. With nonzero κ , the transformed squared-Bessel boundary measure must be weighted back to the intrinsic CIR boundary clock; accumulating unweighted increments in absolute transformed time produces a different clock. The implementation uses local transformed steps and the positive Jacobi bridge weighting, bounding $\kappa\Delta t$ by subdivision. Clock endpoints use the exact affine transition law, while boundary-clock increments replace the conditional bridge local time by its conditional mean. This retains the existing sampled-touch approximation and does not claim to sample the complete local-time distribution.

Next form $\tau_i = t_i + \boldsymbol{\psi}_\tau \mathfrak{t}_i$. In a separate seeded random phase, simulate ζ with exact CIR endpoints over $\Delta\tau_i$, subdividing operational increments when needed. The endpoint-normalized stock return factor of Appendix F applies to this pure operational-time CIR transition, without aggregating different business-time and rough drift coefficients. Its analytic normalizer enforces the discrete stock martingale identity and matches the foreign first state moment. The remaining stock normal is integrated analytically, with a corrected residual variance. Conditional on the initial state and operational step, write $R(\zeta_1)$ for that normalized endpoint factor and $V_{\text{fair}} = \mathbb{E}_d[\int_0^{\Delta\tau} \zeta_u^{-1} du]$. The separate bridge-variance numerical variant uses

$$\begin{aligned} L &= \mathbb{E}_d[-2 \log R(\zeta_1)], & V_{\text{res}} &= V_{\text{fair}} - L = V_{\text{fair}} + 2\mathbb{E}_d[\log R(\zeta_1)], \\ \frac{S_1}{S_0} &= R(\zeta_1) \exp\left(-\frac{1}{2}V_{\text{res}} + \sqrt{V_{\text{res}}}Z\right), & Z &\sim \mathcal{N}(0, 1), \end{aligned} \quad (4.103)$$

where Z is independent of the reciprocal endpoint. The earlier choice $(1 - \rho^2)V_{\text{fair}}$ accounts only for the orthogonal stock Brownian driver and omits fluctuations of the correlated Brownian bridge at a finite step. The correction preserves the discrete martingale identity and foreign first state moment, and enforces the expected log-contract identity

$$\mathbb{E}_d[-2 \log(S_1/S_0)] = V_{\text{fair}}$$

to the accuracy of the fair-variance quadrature. This is expected log-contract variance, not the variance of the terminal log return. It does not reproduce the exact conditional bridge distribution.

For variance–stock correlation $\rho \leq 0$, the normalized exponential tilt with foreign-mean matching satisfies $0 \leq L \leq \rho^2 V_{\text{fair}}$; hence $V_{\text{res}} \geq (1 - \rho^2) V_{\text{fair}} \geq 0$. The corrected numerical variant explicitly restricts pricing to this correlation range. The theoretical dynamics above still permit both signs, subject to their Feller conditions. Exact state endpoints, normalization and this log-contract identity do not make the joint stock transition exact; business-grid and operational-grid refinement remain necessary. Market-delta control variates use fixed market inputs, leave the raw stock forward untouched, and enforce the estimated put–call parity identity.

Independent state drift in the two clocks. The same covariance specification also permits a more flexible drift for ζ . Retain the compulsory Feller-one drift in the diffusion clock, then add independent constant and proportional terms:

$$\mathbb{E}_d[d\zeta] = \frac{1}{2}\boldsymbol{\Psi}_\zeta d\tau + (\alpha - \beta\zeta) dt + (\boldsymbol{\aleph} - \boldsymbol{\beth}\zeta) d\mathfrak{t}, \quad d\tau = dt + \boldsymbol{\Psi}_\tau d\mathfrak{t}. \quad (4.104)$$

Here the compulsory characteristic is fixed at one. The excess constants $\alpha, \boldsymbol{\aleph}$ determine the additional boundary repulsion, independently of the reversion coefficients $\beta, \boldsymbol{\beth}$. Both stock and state quadratic variations still use τ , while the state drift has additional dependence on the separate clocks. The previous model is recovered by setting $\alpha = \frac{1}{2}\boldsymbol{\Psi}_\zeta(\vartheta - 1)$, $\boldsymbol{\aleph} = \boldsymbol{\Psi}_\tau\alpha$, β equal to the previous operational reversion rate, and $\boldsymbol{\beth} = \boldsymbol{\Psi}_\tau\beta$.

The dimensions are $\alpha \sim 1$, $\beta \sim \text{time}^{-1}$, $\boldsymbol{\aleph} \sim \text{time}^\theta$ and $\boldsymbol{\beth} \sim \text{time}^{\theta-1}$. When $\boldsymbol{\Psi}_\tau > 0$, the separate cash Feller characteristics are

$$\vartheta_t = 1 + \frac{2\alpha}{\boldsymbol{\Psi}_\zeta}, \quad \vartheta_{\mathfrak{t}} = 1 + \frac{2\boldsymbol{\aleph}}{\boldsymbol{\Psi}_\zeta\boldsymbol{\Psi}_\tau}. \quad (4.105)$$

The boundary conditions are best enforced without division: the cash excesses $\alpha, \boldsymbol{\aleph}$ and stock-numeraire excesses $\alpha - \rho\sqrt{\boldsymbol{\Psi}_\zeta}$, $\boldsymbol{\aleph} - \rho\sqrt{\boldsymbol{\Psi}_\zeta}\boldsymbol{\Psi}_\tau$ must be nonnegative. For the negative variance–stock correlation used here, nonnegative cash excesses also satisfy the stock conditions. The stock-numeraire map changes these constants and reverses the correlation sign; it leaves the reversion coefficients and clock law unchanged. At $\boldsymbol{\Psi}_\tau = 0$, rough drift may still act on ζ although stock and state innovations use business time.

The Feller-one floor prevents zero attainment. It does not by itself ensure finite inverse-state expectations. The current fair-variance-based simulator therefore requires positive excess drift in each stochastic clock. Stabilizing negative quadratic drift for $1/\zeta$ is stronger again, requiring $\alpha > \boldsymbol{\Psi}_\zeta/2$ and, when the rough diffusion is present, $\boldsymbol{\aleph} > \boldsymbol{\Psi}_\zeta\boldsymbol{\Psi}_\tau/2$.

On a small calendar slab, freeze the ratio of its two clock increments. The aggregated CIR transition then has operational duration $\Delta\tau$ and coefficients

$$\vartheta_{\text{step}} = 1 + \frac{2(\alpha\Delta t + \boldsymbol{\aleph}\Delta\mathfrak{t})}{\boldsymbol{\Psi}_\zeta\Delta\tau}, \quad \beta_{\text{step}} = \frac{\beta\Delta t + \boldsymbol{\beth}\Delta\mathfrak{t}}{\Delta\tau}. \quad (4.106)$$

Subdivide both operational duration and integrated linear drift. Exact endpoints of this frozen CIR transition, normalized stock factors and the bridge-variance correction preserve the discrete stock martingale, foreign first state moment and conditional expected log-contract variance. Aggregation approximates the continuous two-clock drift law and must be checked by calendar-grid refinement. A constant-coefficient whole-interval CIR formula no longer applies in general; the implementation accumulates the local fair variances and audits their paired difference from the conditional log-contract estimate.

Excessively steep long-expiry skew motivates testing a smaller magnitude of negative ρ . This is a calibration hypothesis: the split drift also changes the volatility distribution and hence the skew term structure. Correlation profiles and joint refits are compared using identical market quotes, target weights and numerical settings before attributing an improvement to either mechanism.

Independent smooth and rough reciprocal components. A single stock–variance correlation may not accommodate the observed skew term structure. Assign independent diffusion strengths and correlations to the business and boundary clocks. The business-time coefficient of stock variance remains one, fixing the scale of ζ :

$$\begin{aligned}
\mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)dt, & d\langle\xi\rangle &= 2\xi dt, \\
\mathbb{E}_d[dS] &= 0, & d\langle S\rangle &= \frac{S^2}{\zeta}(dt + \psi_S d\mathfrak{r}), \\
\mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\psi_{\zeta,t} + \alpha - \beta\zeta\right)dt + \left(\frac{1}{2}\psi_{\zeta,\mathfrak{r}} + \mathfrak{N} - \mathfrak{J}\zeta\right)d\mathfrak{r}, & d\langle\zeta\rangle &= \zeta(\psi_{\zeta,t}dt + \psi_{\zeta,\mathfrak{r}}d\mathfrak{r}), \\
d\langle\zeta, S\rangle &= -S(\rho\sqrt{\psi_{\zeta,t}}dt + \mathfrak{N}\sqrt{\psi_{\zeta,\mathfrak{r}}}\psi_S d\mathfrak{r}), & d\langle\xi, S\rangle &= d\langle\xi, \zeta\rangle = 0.
\end{aligned} \tag{4.107}$$

Here ρ and $\mathfrak{N}$ are the smooth and rough variance–stock correlations, opposite in sign to reciprocal-state–stock correlations. Each is in $[-1, 1]$, so each clock’s covariance matrix is positive semidefinite. Constant and proportional drift coefficients remain independent. The compulsory Feller floor is one in each active state innovation clock, without identifying the two actual characteristics. The strengths satisfy $\psi_{\zeta,t} > 0$, $\psi_{\zeta,\mathfrak{r}} \geq 0$, $\psi_S \geq 0$, with dimensions 1, time^θ and time^θ respectively. At $\psi_S = 0$ direct rough stock variance vanishes while rough volatility can remain. At $\psi_{\zeta,\mathfrak{r}} = 0$ rough state innovations vanish without requiring the rough state drift to vanish.

The shared-clock model is the exact subfamily

$$\psi_S = \psi_\tau, \quad \psi_{\zeta,t} = \psi_\zeta, \quad \psi_{\zeta,\mathfrak{r}} = \psi_\zeta\psi_\tau, \quad \mathfrak{N} = \rho. \tag{4.108}$$

Its calibrated parameters thus provide a starting point without losing model expressiveness. Weaker smooth leverage together with stronger rough leverage could

reduce long-expiry skew while retaining short-expiry skew. This is a calibration hypothesis. The correlations cannot be inferred from two skew slopes independently of clock activity and the volatility state distribution.

Cash boundary excesses are $\alpha, \mathbf{N}$. Under the stock numeraire,

$$\alpha_f = \alpha - \rho\sqrt{\Psi_{\zeta,t}}, \quad \mathbf{N}_f = \mathbf{N} - \mathfrak{N}\sqrt{\Psi_{\zeta,\mathfrak{r}}\Psi_S}. \quad (4.109)$$

Both variance–stock correlation signs reverse; reversion, diffusion and clock parameters remain unchanged. Enforce nonnegative excesses in both numeraires without division by rough strengths. For active state diffusions the actual Feller characteristics are $1 + 2\alpha/\Psi_{\zeta,t}$ and $1 + 2\mathbf{N}/\Psi_{\zeta,\mathfrak{r}}$, with the foreign constants from (4.109). Equality gives the nonattainment floor. The fair-variance simulator requires positive excesses in each active state innovation clock for finite inverse-state expectations. Stabilizing negative quadratic drift for $1/\zeta$ is stronger, requiring $\alpha > \Psi_{\zeta,t}/2$ and $\mathbf{N} > \Psi_{\zeta,\mathfrak{r}}/2$ in the corresponding clocks.

Generate the same business-time affine boundary-clock paths first. Freeze the clock ratio within each slab, and define

$$\begin{aligned} \Delta t_\zeta &= \Delta t + \frac{\Psi_{\zeta,\mathfrak{r}}}{\Psi_{\zeta,t}} \Delta \mathfrak{r}, & \Delta \tau &= \Delta t + \Psi_S \Delta \mathfrak{r}, \\ \vartheta_{\text{step}} &= 1 + \frac{2(\alpha \Delta t + \mathbf{N} \Delta \mathfrak{r})}{\Psi_{\zeta,t} \Delta t_\zeta}, & \beta_{\text{step}} &= \frac{\beta \Delta t + \mathfrak{N} \Delta \mathfrak{r}}{\Delta t_\zeta}, \\ \rho_{\text{step}} &= \frac{\rho \sqrt{\Psi_{\zeta,t}} \Delta t + \mathfrak{N} \sqrt{\Psi_{\zeta,\mathfrak{r}} \Psi_S} \Delta \mathfrak{r}}{\sqrt{\Psi_{\zeta,t} \Delta t_\zeta \Delta \tau}}. \end{aligned} \quad (4.110)$$

Cauchy–Schwarz ensures $|\rho_{\text{step}}| \leq 1$. Even when $\rho = \mathfrak{N} = -1$, the effective slab correlation can have smaller magnitude: with both clock increments positive, equality requires $\Psi_{\zeta,\mathfrak{r}}/\Psi_{\zeta,t} = \Psi_S$. This attenuation arises from different smooth and rough diffusion weights; it does not by itself quantify the correlation of returns over a longer horizon. Subdivide both effective durations and bound integrated reversion. Use the resulting noncentral Gamma state transition, with foreign characteristic shift $-2\rho_{\text{step}}\sqrt{\Delta\tau}/(\Psi_{\zeta,t}\Delta t_\zeta)$ for the normalized endpoint stock factor. It matches both the domestic stock martingale and foreign first state moment. Multiply the state-clock fair inverse-variance integral by $\Delta\tau/\Delta t_\zeta$ before subtracting the endpoint’s expected log-contract contribution to obtain residual Gaussian stock variance. The current sampler restricts both leverage correlations to $[-1, 0]$, checks finite second moments of its endpoint tilt, and rejects materially negative residual variance instead of clipping it. These covariance and moment identities hold for the frozen slab law; they do not determine the order of smooth and singular clock events inside a slab. Continuous-model pricing requires business-grid and effective-duration refinement. The nested subfamily retains the earlier sampler’s draw order and prices.

An optional calibration sampler retains the same normal-plus-Gamma end-point distribution but generates its central Gamma component by inverting its distribution function at a fixed uniform draw. This couples changes in the Feller characteristic monotonically, avoiding jumps caused by changing acceptance decisions in a rejection sampler. Separate random keys for each calendar interval and subdivision prevent a local change in subdivision count from shifting all later intervals' random inputs. This changes the numerical coupling, not the transition law or the financial model. Candidates from this search are checked with the original sampler on independent paths before replacing the retained best calibration.

Fitted physical-time rescaling. The holiday-adjusted business clock can be followed by a deterministic increasing map to model time. For calendar expirations T_i , set $u_i = t_{\text{holiday}}(T_i)$ and $u_0 = 0$. One positive speed r_i is assigned to each interval $(u_{i-1}, u_i]$, giving

$$t_{\text{model}}(u) = \sum_{j < i} r_j(u_j - u_{j-1}) + r_i(u - u_{i-1}), \quad u_{i-1} \leq u \leq u_i, \quad r_i > 0. \quad (4.111)$$

The final speed is used for extrapolation. All internal dynamics, including the affine clock-generating state and its boundary clock, run in model time. Market expirations, discount factors, forwards and implied-volatility annualization retain their original calendar convention. This supplies a coherent process across maturities rather than unrelated laws at individual expirations; $r_i = 1$ recovers the fixed business-clock specification exactly.

The first calibration holds the financial parameters fixed and fits all interval speeds to the original quote-error and volume objective, without a smoothness penalty. Report the cumulative map, its interval speeds, the expiry knots and their numerical limits, and the overall ratio $t_{\text{model}}(u_n)/u_n$. Also display the cumulative curve and speeds after dividing by this overall ratio; this separates a common time-unit change from local clock distortions, without changing prices or the fitted map. With one speed per fitted expiry, a smaller in-sample error alone does not establish predictive improvement. The clock is continuous and positive, but its changes of speed are tied to the chosen expiry knots.

If financial parameters are subsequently fitted jointly, fix a dimensional anchor, such as ζ_0 , to remove a redundant global time scale. Indeed, replacing every model time by c times that time leaves the specification equivalent under $\xi_0 \mapsto c\xi_0$, $\zeta_0 \mapsto c\zeta_0$, $\kappa \mapsto \kappa/c$, $\beta \mapsto \beta/c$, $(\boldsymbol{\psi}_S, \boldsymbol{\psi}_{\zeta, \mathfrak{r}}, \mathbf{N}) \mapsto c^\theta(\boldsymbol{\psi}_S, \boldsymbol{\psi}_{\zeta, \mathfrak{r}}, \mathbf{N})$ and $\mathfrak{J} \mapsto c^{\theta-1}\mathfrak{J}$, with the dimensionless parameters unchanged.

An additional Heston stock contribution from the clock driver. The independent reciprocal specification can retain a direct calendar-time stock contri-

bution from ξ , as in the earlier hybrid model. With the normalization of (4.57), introduce $\mathcal{H} \geq 0$ and a separate variance–stock correlation $\rho_\xi \in [-1, 1]$:

$$\begin{aligned}
d\langle S \rangle &= S^2 \left[\left(\frac{\mathcal{H}}{2} \xi + \frac{1}{\zeta} \right) dt + \frac{\boldsymbol{\Psi}_S}{\zeta} d\mathfrak{r} \right], \\
d\langle \xi, S \rangle &= \rho_\xi S \sqrt{\mathcal{H}} \xi dt, \\
d\langle \xi \rangle &= 2\xi dt, \\
d\langle \zeta \rangle &= \zeta (\boldsymbol{\Psi}_{\zeta,t} dt + \boldsymbol{\Psi}_{\zeta,\mathfrak{r}} d\mathfrak{r}), \\
d\langle \zeta, S \rangle &= -S (\rho \sqrt{\boldsymbol{\Psi}_{\zeta,t}} dt + \boldsymbol{\mathfrak{N}} \sqrt{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}} \boldsymbol{\Psi}_S} d\mathfrak{r}), \\
d\langle \xi, \zeta \rangle &= 0.
\end{aligned} \tag{4.112}$$

The domestic drift terms and the definition of the boundary clock are unchanged. The coefficient $\mathcal{H}$ has dimension T^{-2} and is distinct from the Hurst exponent. Setting $\mathcal{H} = 0$ recovers the independent reciprocal model; the two additional parameters do not introduce another state or clock.

The calendar stock martingale is a sum of a Heston contribution driven by ξ and a reciprocal contribution driven by ζ , with independent residual Brownian innovations. Consequently ρ_ξ and ρ are correlations with their respective stock contributions. Their correlations with the total calendar stock innovation are, respectively,

$$\rho_\xi \sqrt{\frac{\mathcal{H}\xi/2}{\mathcal{H}\xi/2 + 1/\zeta}}, \quad -\rho \sqrt{\frac{1/\zeta}{\mathcal{H}\xi/2 + 1/\zeta}}. \tag{4.113}$$

This decomposition makes the covariance matrix positive semidefinite for independent choices of the two component correlations in $[-1, 1]$; there is no additional constraint $\rho_\xi^2 + \rho^2 \leq 1$.

Under the stock numeraire, the clock driver’s drift becomes

$$\mathbb{E}_t[d\xi] = \left[\theta - (\kappa - \rho_\xi \sqrt{\mathcal{H}}) \xi \right] dt. \tag{4.114}$$

The reciprocal foreign drift shifts and its two Feller characteristics remain those specified above. A negative ρ_ξ couples falling stock prices to increases in ξ : its Heston variance contribution increases, while approach to the boundary and rough-clock activity become less likely. The stock-numeraire mean-reversion coefficient is correspondingly larger. This channel can change both the smile and its maturity dependence, rather than only adjusting the maturity’s overall volatility level.

The Heston return factor must be simulated jointly with the retained clock-driver information. An independent Heston return generated after the rough clock would remove the intended leverage dependence. Any finite-step conditional-return approximation must preserve its martingale normalizer and be checked against the

foreign-measure first moment of ξ , as well as covariance and grid convergence. The time-rescaling calibration reported so far uses $\mathcal{H} = 0$ and is retained as the nested comparison, not as a calibration of this extension.

Discrete endpoint return factor. Let $p_\kappa(\xi_1 | \xi_0; \Delta t)$ denote the CIR transition density with drift $\theta - \kappa\xi$ and bracket $2\xi dt$. The separate implementation retains the existing domestic clock paths and uses the following Heston conditional stock factor on each grid interval:

$$M_{\xi,1} = \frac{p_{\kappa-\rho_\xi\sqrt{\mathcal{H}}}(\xi_1 | \xi_0; \Delta t)}{p_\kappa(\xi_1 | \xi_0; \Delta t)}. \quad (4.115)$$

The density ratio integrates to one under the domestic endpoint law, and its first moment is exactly the foreign CIR mean. Products of these factors therefore preserve the discrete martingale condition and the foreign marginal law of the retained ξ endpoints. The existing reciprocal stock factors remain conditionally normalized given the driver and clock paths. The two factors multiply, while their independent Gaussian residual variances add.

For the Heston residual, the first implementation uses the positive trapezoidal area $(\xi_0 + \xi_1)\Delta t/2$ and hence adds variance $(1 - \rho_\xi^2)\mathcal{H}(\xi_0 + \xi_1)\Delta t/4$. This is an approximation to the conditional continuous-time return law: it omits unresolved fluctuations of the integrated driver and of its bridge-conditioned boundary clock. Exact discrete normalization must not be confused with exact continuous-time option pricing. Its numerical evaluation uses scaled Bessel functions and the central Gamma limit, with no empirical stock rescaling. The $\mathcal{H} = 0$ branch reuses the original snapshots and random draws exactly.

Conditional VIX approximations for $\mathcal{H} = 0$

For the independent reciprocal model, a conditional VIX approximation must include the total stock quadratic variation in (4.107). Let T denote calendar time, $\Delta_c = 30/365$, and

$$\begin{aligned} t(T) &= t_{\text{model}}(t_{\text{holiday}}(T)), \\ U_T &= t(T + \Delta_c) - t(T). \end{aligned} \quad (4.116)$$

Both the affine clock driver and the reciprocal state evolve over the internal horizon U_T . The annualization denominator remains the calendar window Δ_c . Conditional

on the natural state $(\xi_{t(T)}, \zeta_{t(T)}) = (\xi, \zeta)$, define

$$\begin{aligned} V_T(\xi, \zeta) &= \mathbb{E}_d \left[\int_0^{U_T} \frac{du + \boldsymbol{\Psi}_S d\mathfrak{r}_{t(T)+u}}{\zeta_{t(T)+u}} \middle| \xi, \zeta \right], \\ \text{VIX}_T^{\text{cont}}(\xi, \zeta) &= 100 \sqrt{\frac{V_T(\xi, \zeta)}{\Delta_c}}. \end{aligned} \quad (4.117)$$

Here the superscript denotes the continuous-model variance target. For a continuous discounted stock with finite conditional expected quadratic variation, Itô's formula also gives

$$V_T(\xi, \zeta) = -2\mathbb{E}_d \left[\log \frac{S_{t(T)+U_T}}{S_{t(T)}} \middle| \xi, \zeta \right]. \quad (4.118)$$

This is expected log-contract variance, not the variance of the terminal log return. Neither $(1 - \rho^2)$ nor $(1 - \mathfrak{p}^2)$ multiplies the total variance: those factors concern the orthogonal return decomposition. The market VIX methodology uses selected option strikes and maturity interpolation, so its finite option portfolio need not equal this continuous-model target exactly.³

The canonical conditional mean clock. Write $p = 1 - \theta$, and let

$$F_\kappa(u) = \frac{e^{\kappa u} - 1}{\kappa}, \quad D_\kappa(u) = \frac{1 - e^{-\kappa u}}{\kappa}, \quad (4.119)$$

with both expressions interpreted as u at $\kappa = 0$. For the canonical physical CIR boundary clock, the exact conditional mean and its rate are

$$\begin{aligned} m(u \mid \xi) &= \mathbb{E}_d [\mathfrak{r}_{t+u} - \mathfrak{r}_t \mid \xi_t = \xi] = \int_0^u r(v \mid \xi) dv, \\ r(u \mid \xi) &= \frac{p}{\Gamma(\theta)} D_\kappa(u)^{-\theta} \exp \left(-\frac{\xi}{F_\kappa(u)} \right), \quad u > 0. \end{aligned} \quad (4.120)$$

This mean is time homogeneous in internal model time; no absolute-time exponential factor is present. At $\kappa = 0$ and $\xi = 0$ it reduces to $m(u \mid 0) = u^p / \Gamma(\theta)$. With $\kappa > 0$, using $F_\kappa(u)^p / \Gamma(\theta)$ at the boundary instead would describe the unweighted transformed clock. The integrable $u^{-\theta}$ singularity at $\xi = 0$ is handled through clock increments or a suitable quadrature coordinate, rather than evaluation of $r(0 \mid 0)$.

³Cboe's official description of the SPX volatility-index methodology specifies a 30-day target and calendar annualization, with finite strike selection and two-term interpolation; the VIXMO variant uses the same variance calculation: https://cdn.cboe.com/api/global/us_indices/governance/Volatility_Index_methodology_VIXMO_only.pdf.

A frozen conditional mean-rate approximation. For $U = U_T > 0$, replace the future rough-clock rate by the constant $\bar{r} = m(U \mid \xi)/U$, and define

$$\begin{aligned} q &= \Psi_{\zeta,t} + \Psi_{\zeta,\bar{r}}, & a &= \alpha + \aleph \bar{r}, \\ b &= \beta + \beth \bar{r}, & k &= 1 + \Psi_S \bar{r}, & \nu &= 1 + \frac{2a}{q}. \end{aligned} \quad (4.121)$$

The reciprocal state then has constant affine drift $q/2 + a - b\zeta$ and bracket $q\zeta du$. Retaining its diffusion gives the fixed-CIR fair-variance formula of (F.19):

$$\widehat{V}_T^{\text{frozen}}(\xi, \zeta) = k y \left(\frac{F_b(U)}{\zeta}; \nu, q \right), \quad q > 0, \quad \nu > 1. \quad (4.122)$$

It is exact for this constant clock ratio. It does not replace the reciprocal state by its conditional mean; such a replacement would add an inverse-state approximation.

The same exact fair-variance formula is derived in the companion *VIX and CDS in Reciprocal Affine Stochastic Vol Model*, equations (3.1) and (4.2), both on page 4.⁴ Its parameters $(\xi, \theta, \gamma, \mu)$ correspond here to $(\zeta, q/2 + a, q, b)$. In the pure smooth model this becomes $(\zeta, \Psi_{\zeta,t}/2 + \alpha, \Psi_{\zeta,t}, \beta)$ and $x = F_\beta(U)/\zeta$ exactly, without a mean-state approximation. For $\chi = 2/(qx)$ and $\delta = \nu - 1 > 0$, the corrected equation (3.1) can be written stably as

$$\begin{aligned} \frac{\partial y}{\partial x} &= \chi \int_0^1 (1-v)^{\delta-1} e^{-xv} dv = \frac{\chi}{\delta} {}_1F_1(1; \nu; -\chi), \\ y(x; \nu, q) &= \frac{2}{q\delta} \left[E_1(\chi) + \int_0^\chi \frac{1 - (1-v/\chi)^\delta}{v} e^{-v} dv \right], \end{aligned} \quad (4.123)$$

where $E_1(\chi) = \int_\chi^\infty e^{-v} dv/v$. Differentiation of the second line gives the first. They satisfy the paper's fair-variance ODE with $y(0) = 0$ and $y'(0) = 1$. Since that ODE is singular at zero, the fair solution also selects logarithmic growth at infinity, excluding the additional homogeneous solution. Thus the smooth and frozen-rate formulas are exact special-function or integral representations; numerical positive quadrature evaluates them, rather than approximating their reciprocal diffusion dynamics.

The zero-diffusion limit holds the immigration a fixed and is

$$\widehat{V}_T^{\text{frozen}} = \begin{cases} \frac{k}{a} \log \left(1 + \frac{a F_b(U)}{\zeta} \right), & a > 0, \\ \frac{k F_b(U)}{\zeta}, & a = 0. \end{cases} \quad (4.124)$$

Holding a finite ν fixed while $q \downarrow 0$ instead forces $a = q(\nu - 1)/2$ to vanish.

⁴Gregory Pelts, April 3, 2015; source in the companion project `Reciprocal Affine Model/main.tex`.

Retaining the full conditional mean-clock path. A finer approximation replaces $d\mathfrak{r}_{t+u}$ by $dm(u \mid \xi)$ throughout the window, with the entire mean path conditioned on its initial natural state ξ . Its reciprocal state remains stochastic:

$$\begin{aligned}\mathbb{E}_d[d\zeta_u] &= \left[\frac{q(u)}{2} + \alpha + \mathfrak{N}r(u \mid \xi) - (\beta + \mathfrak{I}r(u \mid \xi))\zeta_u \right] du, \\ d\langle \zeta \rangle_u &= q(u)\zeta_u du, \quad q(u) = \mathfrak{V}_{\zeta,t} + \mathfrak{V}_{\zeta,\mathfrak{r}}r(u \mid \xi), \\ \widehat{V}_T^{\text{path}} &= \int_0^U (1 + \mathfrak{V}_{Sr}(u \mid \xi)) \mathbb{E}_d[\zeta_u^{-1}] du.\end{aligned}\tag{4.125}$$

In general this time-inhomogeneous CIR expectation cannot be reduced to one evaluation of y using only $m(U \mid \xi)$.

For a partition $0 = u_0 < \dots < u_n = U$, freeze the clock ratio on each slab using $\Delta u_i = u_i - u_{i-1}$ and $\Delta m_i = m(u_i \mid \xi) - m(u_{i-1} \mid \xi)$. Set

$$\begin{aligned}Q_i &= \mathfrak{V}_{\zeta,t}\Delta u_i + \mathfrak{V}_{\zeta,\mathfrak{r}}\Delta m_i, & e_i &= \alpha\Delta u_i + \mathfrak{N}\Delta m_i, \\ R_i &= \beta\Delta u_i + \mathfrak{I}\Delta m_i, & \pi_i &= \Delta u_i + \mathfrak{V}_S\Delta m_i, \\ \nu_i &= 1 + \frac{2e_i}{Q_i}, & g_i &= \frac{Q_i}{2} \frac{1 - e^{-R_i}}{R_i}, & q_i &= e^{-R_i}.\end{aligned}\tag{4.126}$$

The ratio in g_i has limit one at $R_i = 0$. If $L_i(\lambda) = \mathbb{E}_d[e^{-\lambda\zeta_{u_i}}]$, exact affine transitions of this frozen-slab model give

$$L_0(\lambda) = e^{-\lambda\zeta}, \quad L_i(\lambda) = (1 + g_i\lambda)^{-\nu_i} L_{i-1} \left(\frac{q_i\lambda}{1 + g_i\lambda} \right).\tag{4.127}$$

This recursion retains the complete reciprocal-state marginal distribution. For $\delta_i = \nu_i - 1 > 0$ and $\eta_i = q_i/g_i$, the expected total stock variance on slab i has the positive representation

$$\begin{aligned}W_i &= \frac{2\pi_i}{Q_i} \left[\frac{1}{\delta_i} \int_{\eta_i}^{\infty} \frac{L_{i-1}(\lambda)}{\lambda} d\lambda \right. \\ &\quad \left. + \int_0^{\eta_i} \frac{1 - (1 - \lambda/\eta_i)^{\delta_i}}{\delta_i} \frac{L_{i-1}(\lambda)}{\lambda} d\lambda \right].\end{aligned}\tag{4.128}$$

Thus $\sum_i W_i$ approximates (4.125) without sampling reciprocal states. Positive quadrature in a scaled or logarithmic Laplace coordinate resolves narrow small-slab peaks. Clock and reciprocal-state diffusion have independent coefficients, so this distribution propagation cannot be replaced by advancing only $\mathbb{E}[\zeta]$.

The compulsory descriptor $\vartheta = 1$ is only the nonattainment floor. The actual cash characteristics are $1 + 2\alpha/\mathfrak{V}_{\zeta,t}$ and $1 + 2\mathfrak{N}/\mathfrak{V}_{\zeta,\mathfrak{r}}$ in their active innovation clocks.

Strictly positive excesses ensure $\nu_i > 1$ and finite first inverse-state expectations in the frozen CIR slabs. At characteristic one, their fair variance diverges; no finite value is obtained by clipping that characteristic. Characteristics close to one require quadrature checks. Divergence of an inverse-endpoint second moment at characteristic at most two does not by itself establish divergence of the continuously integrated inverse-area second moment.

Refine both the slab partition and the fair-variance quadrature, and compare conditional total variance before taking square roots. Sampled canonical clock paths provide a benchmark for replacing their distribution by m , with reciprocal diffusion retained. Sampling uncertainty is reported separately from clock-bridge and interval approximations. Neither mean-clock approximation has a universal bias sign, and agreement at the initial state does not establish agreement across option-expiry states.

An atom-preserving clock-distribution closure. The conditional mean alone discards the intermittency of the boundary clock. Let $A_u = \mathfrak{r}_{t+u} - \mathfrak{r}_t$ and $M_n(u \mid \xi) = \mathbb{E}_d[A_u^n \mid \xi_t = \xi]$. The strong Markov property at boundary-clock increments gives the exact Kac moment recursion

$$M_0(u \mid \xi) = 1, \quad M_n(u \mid \xi) = n \int_0^u r(v \mid \xi) M_{n-1}(u - v \mid 0) dv. \quad (4.129)$$

In particular $M_1 = m$ and $M_2(u \mid \xi) = 2 \int_0^u r(v \mid \xi) m(u - v \mid 0) dv$, using the physical rate in (4.120) also when $\kappa > 0$. At $\kappa = 0$ and $\xi = 0$, the normalization can be checked from

$$M_n(u \mid 0) = \frac{n!}{\Gamma(1 + np)} \left(\frac{\Gamma(1 + p)}{\Gamma(\theta)} u^p \right)^n. \quad (4.130)$$

The exact no-contact atom is

$$p_0(u \mid \xi) = \mathbb{P}_d(A_u = 0 \mid \xi_t = \xi) = \frac{1}{\Gamma(p)} \int_0^{\xi/F_\kappa(u)} v^{p-1} e^{-v} dv. \quad (4.131)$$

Thus $p_0(u \mid 0) = 0$ for $u > 0$.

A computable marginal closure preserves this atom and fits a gamma law to the positive clock. With $h = 1 - p_0$, put

$$\mu_+ = \frac{M_1}{h}, \quad v_+ = \frac{M_2}{h} - \mu_+^2, \quad A_u \approx p_0 \delta_0 + h \text{ Gamma} \left(\frac{\mu_+^2}{v_+}, \frac{v_+}{\mu_+} \right). \quad (4.132)$$

The gamma arguments denote shape and scale. If $v_+ = 0$, the positive branch is a point mass at μ_+ . This is a two-moment closure; higher moments and tails are not identified by these formulas.

A fast total-clock mixture evaluates (4.122) at random constant ratios A_U/U and averages with the law in (4.132). It retains the reciprocal-state diffusion but discards within-window clock timing. A path version uses the gamma-mixture quantile Q_u at each partition time and a common uniform variable Z :

$$\hat{A}_{u_i} = Q_{u_i}(Z), \quad Z \sim \text{Uniform}(0, 1). \quad (4.133)$$

On each increasing quantile path, replace Δm_i in (4.126) by $\hat{A}_{u_i} - \hat{A}_{u_{i-1}}$, propagate (4.127), and integrate (4.128). Positive quadrature then averages these conditional fair areas. The complete reciprocal-state law is propagated on every path. Gamma quadrature can reproduce the first two terminal clock moments exactly while earlier contact-probability cutoffs still require refinement. Report quadrature errors in the mean, second moment, and zero atom across the entire time partition.

The shared quantile imposes a comonotonic coupling, not the actual joint law of the CIR boundary clock. It can overstate temporal covariance and smooth the random contact episodes. Smooth and rough reciprocal generators generally do not commute, so matching terminal clock moments, or even all one-time marginals, does not identify expected stock variance. This limitation is especially material when the rough characteristic is close to one. Finer time partitions and more positive quadrature nodes control numerical errors of the closure; they do not remove its joint-law error. Sampled-clock comparisons must also refine their clock bridges and within-interval ordering before serving as a continuum benchmark.

A common per-expiry VIX forward scale. For VIX option expiry T , let $Y_T^{(j)}$ be simulated VIX levels obtained by applying the conditional variance approximation to each expiry state. A common forward alignment uses

$$\bar{Y}_T = \frac{1}{N} \sum_{j=1}^N Y_T^{(j)}, \quad a_T = \frac{F_T^{\text{target}}}{\bar{Y}_T}, \quad \tilde{Y}_T^{(j)} = a_T Y_T^{(j)}. \quad (4.134)$$

Use the same a_T for calls and puts of that expiry. With discount factor D_T , the resulting sample prices satisfy

$$\hat{C}_T(K) - \hat{P}_T(K) = D_T(F_T^{\text{target}} - K). \quad (4.135)$$

An observed VIX futures quote supplies the preferred target. When that quote is absent from the retained cache, the fallback is an explicitly identified option-implied forward from matched call and put quotes, $F_T^{\text{opt}} = K + (C_T^{\text{mkt}} - P_T^{\text{mkt}})/D_T$. Report the raw Monte Carlo mean $\bar{Y}_T$, the target and its source, and the scale a_T for every expiry. This aligns the mean of VIX levels; averaging conditional variance before taking its square root gives a different quantity.

The scale can absorb differences between the continuous variance target and the market's finite option portfolio, including strike and tail cutoffs. It can also absorb model or approximation error. Forward alignment therefore does not validate the conditional variance approximation or remove errors in the distribution's shape; an option-implied target estimated from the same quotes is not an independent validation input.

Jump-free reciprocal model with a common Feller characteristic

The independent-clock specification admits the following constrained version: the smooth and rough reciprocal innovations have the same actual cash-numeraire Feller characteristic $\vartheta > 1$. The affine driver's characteristic $\theta \in (0, 1)$ remains independent of ϑ . Here t is internal business/model time, obtained from the deterministic physical-time map; there is no additional random observation clock. For the initial specification set $\mathcal{H} = 0$. Its complete drift and covariance equations are

$$\begin{aligned} \mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)dt, & d\langle\xi\rangle &= 2\xi dt, \\ \mathbb{E}_d[dS] &= 0, & d\langle S\rangle &= \frac{S^2}{\zeta}(dt + \boldsymbol{\psi}_S d\mathfrak{r}), \\ \mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\vartheta\boldsymbol{\psi}_{\zeta,t} - \beta\zeta\right)dt + \left(\frac{1}{2}\vartheta\boldsymbol{\psi}_{\zeta,\mathfrak{r}} - \boldsymbol{\beth}\zeta\right)d\mathfrak{r}, & d\langle\zeta\rangle &= \zeta(\boldsymbol{\psi}_{\zeta,t}dt + \boldsymbol{\psi}_{\zeta,\mathfrak{r}}d\mathfrak{r}), \\ d\langle\zeta, S\rangle &= -S(\rho\sqrt{\boldsymbol{\psi}_{\zeta,t}}dt + \boldsymbol{\mathfrak{N}}\sqrt{\boldsymbol{\psi}_{\zeta,\mathfrak{r}}}\boldsymbol{\psi}_S d\mathfrak{r}), & d\langle\xi, S\rangle &= d\langle\xi, \zeta\rangle = 0. \end{aligned} \tag{4.136}$$

The coefficients satisfy $\kappa, \beta, \boldsymbol{\beth} \geq 0$, $\xi_0 \geq 0$, $\zeta_0 > 0$, $\boldsymbol{\psi}_{\zeta,t} > 0$, $\boldsymbol{\psi}_{\zeta,\mathfrak{r}}, \boldsymbol{\psi}_S \geq 0$ and $|\rho|, |\boldsymbol{\mathfrak{N}}| \leq 1$. The current bridge-return sampler further restricts $\rho, \boldsymbol{\mathfrak{N}} \leq 0$. As before, these are variance-stock correlations, with the opposite sign for reciprocal-state-stock correlations. The deterministic smooth stock coefficient is one, fixing the reciprocal-state scale. The optional Heston contribution is precisely (4.112); it is not required for this constraint or for the variance table.

The old floor-plus-excess-drift representation gives exactly the same parameters after the substitution

$$\alpha = \frac{\vartheta - 1}{2}\boldsymbol{\psi}_{\zeta,t}, \quad \boldsymbol{\mathfrak{N}} = \frac{\vartheta - 1}{2}\boldsymbol{\psi}_{\zeta,\mathfrak{r}}. \tag{4.137}$$

Thus α and $\boldsymbol{\mathfrak{N}}$ are derived quantities in this version, rather than additional free calibration variables. The smooth and rough reversion rates, diffusion strengths and leverage coefficients remain independent. Their units remain those of the independent-clock model: $[\xi] = [\zeta] = T$, $[\mathfrak{r}] = T^{1-\theta}$, $[\kappa] = [\beta] = T^{-1}$, $[\boldsymbol{\beth}] = T^{\theta-1}$, $[\boldsymbol{\psi}_{\zeta,t}] = 1$ and $[\boldsymbol{\psi}_{\zeta,\mathfrak{r}}] = [\boldsymbol{\psi}_S] = T^\theta$; $\theta, \vartheta, \rho, \boldsymbol{\mathfrak{N}}$ are dimensionless.

In particular, the effective domestic characteristic in every frozen smooth/rough slab is independent of its sampled clock ratio:

$$1 + \frac{2(\alpha\Delta t + \mathfrak{N}\Delta\mathfrak{r})}{\mathfrak{V}_{\zeta,t}\Delta t + \mathfrak{V}_{\zeta,\mathfrak{r}}\Delta\mathfrak{r}} = \vartheta. \quad (4.138)$$

This removes the mixture of different domestic characteristics in (4.110). It does not remove the approximation of the smooth/rough event ordering inside a finite slab. The new implementation translates the descriptor and reuses the existing transition, stock normalizer and random-input law; the unconstrained and jump-clock implementations and their results are retained separately. A fit with unequal characteristics must be projected onto a chosen common ϑ and recalibrated. Its original fit error is not automatically attainable under this additional constraint.

Numeraire conditions and inverse-state integrability. For an active rough state diffusion, the two stock-numeraire characteristics become

$$\vartheta_{\mathfrak{f},t} = \vartheta - \frac{2\rho}{\sqrt{\mathfrak{V}_{\zeta,t}}}, \quad \vartheta_{\mathfrak{f},\mathfrak{r}} = \vartheta - 2\mathfrak{N}\sqrt{\frac{\mathfrak{V}_S}{\mathfrak{V}_{\zeta,\mathfrak{r}}}}. \quad (4.139)$$

Both must satisfy their nonattainment floors, individually. Equal cash characteristics do not require equal stock characteristics: the numeraire shifts depend on the two independent leverage channels. The rough condition is omitted when $\mathfrak{V}_{\zeta,\mathfrak{r}} = 0$; it must not be evaluated by division by a zero diffusion strength. In the implemented nonpositive-correlation regime both stock characteristics are at least the common cash characteristic.

For any positive elapsed CIR time, the inverse-state mean is infinite at cash characteristic one, even though that value is the nonattainment floor. Consequently its expected fair variance is infinite at that critical characteristic. A strictly larger ϑ makes the first inverse-state expectation finite, but approaching one creates a poorly conditioned fair-variance calculation and strong sensitivity to rare small reciprocal states. Finite results on a coarse grid do not establish a finite critical-characteristic limit. The threshold two instead concerns the instantaneous second inverse-state moment and the sign of the inverse state's quadratic drift; it is not a general criterion for every time-integrated variance moment. Calibration bounds and variance-table convergence checks therefore state their distance from the critical first-moment floor explicitly.

Mean reversion and the conditional rough clock. The exact nonzero- κ conditional clock mean is already (4.120); it is not obtained by inserting an

unweighted transformed time into the zero-reversion formula. For comparison, the first-contact probability over $U > 0$ is

$$P_\kappa(U \mid \xi) = Q\left(1 - \theta, \frac{\xi}{F_\kappa(U)}\right), \quad \widehat{m}_\kappa(U \mid \xi) = m_0(U \mid \xi) \frac{P_\kappa(U \mid \xi)}{P_0(U \mid \xi)}, \quad (4.140)$$

where Q is the regularized upper incomplete Gamma function and F_κ is defined in (4.119). The probability ratio corrects contact frequency while retaining the zero-reversion amount of clock conditional on contact. It can improve the zero-reversion estimate, but is still an approximation. At $\xi = 0$ both contact probabilities are one, so this ratio misses the increase in the physical boundary-clock mean caused by mean reversion. The exact integral remains the default mean-clock input and the diagnostic reference. The optional ratio is evaluated through the conditional positive-contact mean when necessary to avoid division of underflowed probabilities.

Variance-only tabulation. Calibrate the jump-free specification with the common characteristic constraint first. Check the selected parameters using fresh path batches and a refined outer time grid, then freeze the complete parameter vector and fitted physical-time map. Only then build the conditional variance tables. A parameter change during subsequent calibration invalidates those tables and requires retabulation; a table from an earlier candidate is not used to certify the final calibrated model. For a calendar expiry T , compose the holiday and fitted physical-time maps into $u(T)$ and use the future internal duration. Write Δ_T for the future variance window in ACT/365 calendar years; the standard monthly window has $\Delta_T = 30/365$:

$$U_T = u(T + \Delta_T) - u(T), \quad V_T(\xi, \zeta) = \mathbb{E}_d \left[\int_0^{U_T} \frac{ds + \boldsymbol{\psi}_S d\mathfrak{r}_{u(T)+s}}{\zeta_{u(T)+s}} \middle| \xi, \zeta \right]. \quad (4.141)$$

Tabulate this conditional integrated variance on a grid of ξ and ζ , including boundary starts and the states observed at fitted option expiries. The default table generator samples only the clock driver and its boundary-clock increments, integrating out the reciprocal innovations. A sampled reciprocal-state alternative provides a separate check. Neither method generates an underlying level, stock endpoint tilt or idiosyncratic stock shock.

On each sampled clock slab, write

$$\begin{aligned} \Delta t_\zeta &= \Delta s + \frac{\boldsymbol{\psi}_{\zeta, \mathfrak{r}}}{\boldsymbol{\psi}_{\zeta, t}} \Delta \mathfrak{r}, & \Delta \tau &= \Delta s + \boldsymbol{\psi}_S \Delta \mathfrak{r}, \\ \beta_{\text{step}} &= \frac{\beta \Delta s + \boldsymbol{\mathfrak{r}} \Delta \mathfrak{r}}{\Delta t_\zeta}, & W_{\text{step}}(\zeta) &= \frac{\Delta \tau}{\Delta t_\zeta} y \left(\frac{F_{\beta_{\text{step}}}(\Delta t_\zeta)}{\zeta}; \vartheta, \boldsymbol{\psi}_{\zeta, t} \right). \end{aligned} \quad (4.142)$$

The function y is the mean-reverting smooth reciprocal fair-variance integral in (4.123). The sampled-state alternative advances the exact frozen-slab noncentral-Gamma endpoint and accumulates W_{step} conditional on its starting state, integrating out its unresolved reciprocal bridge.

The common characteristic also permits integrating out all reciprocal state innovations, conditional on each sampled rough-clock path. Set $R_i = \beta \Delta s_i + \mathfrak{z} \Delta \mathfrak{r}_i$, $G_0 = 0$ and $D_i = \exp(-\sum_{j \leq i} R_j)$. Its endpoint marginals satisfy

$$G_i = e^{-R_i} G_{i-1} + \frac{\mathfrak{v}_{\zeta,t} \Delta t_{\zeta,i}}{2} \frac{1 - e^{-R_i}}{R_i},$$

$$\zeta_i \mid \mathfrak{r} \stackrel{d}{=} G_i \Gamma\left(\vartheta + \text{Poisson}\left(\frac{\zeta D_i}{G_i}\right), 1\right).$$

The ratio at $R_i = 0$ is interpreted as one. Evaluate the marginal log moments $\ell_i = \mathbb{E}[\log \zeta_i \mid \mathfrak{r}]$ by positive noncentral-Gamma quadrature, with $\ell_0 = \log \zeta$. Itô's formula for $\log \zeta$ then gives the conditional slab area without a reciprocal-state draw:

$$\mathbb{E}[W_i \mid \mathfrak{r}] = \frac{\Delta \tau_i}{\Delta t_{\zeta,i}} \frac{\ell_i - \ell_{i-1} + R_i}{\frac{1}{2} \mathfrak{v}_{\zeta,t} (\vartheta - 1)}.$$

Average the summed areas over the sampled clock paths. This removes reciprocal-state sampling error, while retaining rough-clock randomness and the finite-grid clock-ordering approximation.

The table records the parameter vector, future internal duration and time grid, path counts, random seeds, quadrature settings, sampled clock mean and exact conditional clock mean. Refine the outer clock partition, especially its first intervals near $\xi = 0$, as well as the reciprocal substeps and fair-variance quadrature. Increasing reciprocal subdivision alone cannot recover unresolved rough-clock ordering. Use independent path batches to report sampling error separately from refinement changes. Interpolate positive integrated variances linearly in ξ . The basic method is linear in $\log \zeta$; an optional shape-preserving cubic interpolation of $\log V$ against $\log \zeta$ retains positivity and reduces the positive interpolation bias in the large- ζ inverse-state asymptote. Check both methods against independent off-grid estimates. Identify states outside the table and extend the grid instead of silently extrapolating or clipping them. A table for one future internal duration must not be substituted for a different expiry's duration without an explicit additional interpolation or a separate table.

Paired fair-variance and forward comparison. The numerical calibration imposes an explicit margin from the critical floor, $\vartheta \geq 1.5$, for both cash characteristics. This is a chosen calibration restriction, not an additional nonattainment theorem. After freezing the fit, generate one common bank of 2^{20} variance

paths, shared across all VIX expiries. The calendar partition contains every expiry and future-window endpoint and has steps no longer than one quarter of a day. Simulate ξ , its boundary clock and ζ , and accumulate

$$A_T = \int_{u(T)}^{u(T+\Delta_T)} \frac{ds + \boldsymbol{\psi}_S d\mathfrak{r}_s}{\zeta_s}.$$

No stock level or stock-correlation tilt is needed. The full-path reference samples reciprocal midpoint states and integrates their inverse values, rather than substituting a conditional fair-area formula. Finer reciprocal subdivision and independent batch diagnostics check the integration and sensitivity to rare small states separately from the outer clock partition.

Retain each path's expiry state and A_T . Build the conditional estimator $\widehat{V}_T(\xi_T, \zeta_T)$ over the full observed state range, without clipping tail states. Its future partition is the corresponding segment of the reference partition, mapped as $u(T+s) - u(T)$. Then evaluate that estimator on the same retained paths; additional future-clock draws used to construct the table do not replace this common outer sample. Near the boundary, interpolate in $\xi^{1-\theta}$ to resolve the leading boundary-clock cusp. Compare it with ordinary ξ interpolation on the same table using independent conditional estimates at the same off-grid states; neither comparison supplies a fitted level correction. With bars denoting averages over those paths, report four levels per expiry, in VIX points:

$$100\sqrt{\frac{360}{365\Delta_T}} \overline{A_T}, \quad 100\sqrt{\frac{360}{365\Delta_T}} \widehat{\overline{V_T}}, \quad 100\sqrt{\frac{360}{365\Delta_T}} \widehat{\widehat{V_T}}, \quad F_T^{(360)}. \quad (4.143)$$

The first two agree in the exact conditional-expectation and converged simulation limit, by the tower property. Their paired difference measures estimator and integration error; the difference between the second and third displays the Jensen effect. In the absence of observed futures, the final column is explicitly identified as an option-parity synthetic forward. No fitted multiplier matches any model column to it.

This comparison deliberately uses a 360-day annualization year, so the standard monthly factor is 360/30. The official Cboe convention uses a 365-day year; retain the original market forward and display $F_T^{(360)} = \sqrt{360/365} F_T^{(365)}$ for consistent units. Weekly SOQ windows use their nominal 30 days plus 390 minutes in place of 30 days.⁵ Record the settlement and nominal future-window conventions, and flag future windows that extrapolate beyond the fitted physical-time knots.

⁵Cboe, *VIX Methodology*, https://cdn.cboe.com/resources/vix/VIX_Methodology.pdf; the nominal weekly SOQ window is described in <https://www.cboe.com/tradable-products/vix/faqs/>.

Future averaging and the reciprocal tail. An instantaneous elasticity above one does not determine the sign of the VIX option smile. Even in the pure smooth reciprocal case, the mean-reverting fair-variance integral is

$$W(U \mid \zeta) = \frac{\mathbb{E}[\log \zeta_U \mid \zeta] - \log \zeta + \beta U}{\frac{1}{2} \boldsymbol{\Psi}_{\zeta,t}(\vartheta - 1)}.$$

For fixed $U > 0$ set $G_U = \frac{1}{2} \boldsymbol{\Psi}_{\zeta,t} D_\beta(U)$, with $D_\beta(U) = (1 - e^{-\beta U})/\beta$ and its continuous zero-rate limit. As the initial reciprocal state decreases to zero, the noncentrality vanishes and $\mathbb{E}[\log \zeta_U] \rightarrow \log G_U + \psi(\vartheta)$. Consequently,

$$W(U \mid \zeta) = \frac{\log(G_U/\zeta) + \psi(\vartheta) + \beta U + o(1)}{\frac{1}{2} \boldsymbol{\Psi}_{\zeta,t}(\vartheta - 1)}.$$

The corresponding conditional VIX level grows only as the square root of a logarithm of the initial variance, whereas the instantaneous smooth volatility grows as $\zeta^{-1/2}$. Thus future averaging can substantially compress the right tail even when typical levels require only a modest forward correction. This identifies a mechanism, not a general sign theorem for the full model's VIX implied skew. Compare the tabulated full future expectation, the pure smooth future expectation, and the instantaneous smooth-volatility proxy on the same sampled expiry states; keep their respective forward-normalization factors explicit. Only the first is the model's VIX forecast.

VIX levels, skew and parity. Evaluate the table at each simulated option-expiry state and set

$$Y_T = 100 \sqrt{\frac{V_T(\xi_{u(T)}, \zeta_{u(T)})}{30/365}}, \quad a_T = \frac{F_T^{\text{target}}}{N^{-1} \sum_{j=1}^N Y_T^{(j)}}, \quad \tilde{Y}_T^{(j)} = a_T Y_T^{(j)}. \quad (4.144)$$

The denominator is calendar time, regardless of the internal dynamics' time rescaling. Use the same positive scaling factor for calls and puts, so their discounted Monte Carlo prices satisfy $C_T(K) - P_T(K) = D_T(F_T^{\text{target}} - K)$ exactly on the sample. Report the scale and the target's source; when a futures quote is absent, an identified matched-call/put implied forward can supply the target. Also report the unscaled model mean and its difference from that target. Forward matching enforces parity; it does not validate the conditional fair-variance calculation. In particular, omitted positive far-strike option contributions reduce a truncated replication sum. Such truncation cannot explain a full-model estimate lying below its own otherwise identical truncated replication estimate. A shortfall relative to market therefore requires checking numerical variance and clock approximations,

the fitted state distribution, and model consistency with the option prices; it must not be attributed automatically to a portfolio cutoff. Plot market and model Black implied volatilities against $\log(K/F_T^{\text{target}})$ to compare smile shape after forward alignment. The reciprocal variance state's 3/2 diffusion elasticity alone does not settle this conditional thirty-day variance smile. Check its shape using the tabulated full conditional variance, with state-range and time-grid refinement, before drawing a conclusion from a mean-clock approximation's skew.

Sampling the boundary-clock bridge. The conditional bridge mean used in the original clock sampler preserves the first clock moment up to quadrature error, but suppresses higher moments. Matching the exact conditional mean therefore does not establish an accurate future-variance distribution. An optional reference sampler retains exactly the same affine endpoints and first/last contact draws and restores the independent clock amount between these contacts.

Put $p = 1 - \theta$ and

$$A(u) = \frac{\sin u}{\sin(pu)^p \sin((1-p)u)^{1-p}}, \quad M = p^{-p}(1-p)^{-(1-p)}.$$

Draw U on $(0, \pi)$ with density proportional to $A(u)$, by accepting a uniform proposal with probability $A(u)/M$, and independently draw $G \sim \Gamma(2-p, 1)$. Multiply the conditional bridge mean by

$$R = \frac{\Gamma(2p)}{\Gamma(p)} A(U) G^{1-p}, \quad \mathbb{E}[R^n] = \frac{n! \Gamma(2p)^n \Gamma(p)^{1-n}}{\Gamma((n+1)p)}.$$

In particular $\mathbb{E}[R] = 1$. This is the size-biased inverse-positive- stable bridge law; at $\kappa = 0$ it samples the exact clock amount conditional on the contact times. A separate random stream preserves the original affine endpoint inputs and draws the multiplier in every lane, including lanes without contact.

For positive κ , keeping the original averaged within-bridge weight leaves a controlled approximation. Coupling with the exact weighted bridge gives a relative weight error bounded by $\exp(p\kappa\Delta t) - 1$ on an outer slab. Report its maximum across the actual grid. Apply the multiplier to each individual contact slab, never to an aggregate containing hidden clock substeps. This improvement restores clock randomness but does not remove the finite-slab smooth/rough event- ordering approximation. Recheck the stock calibration with this sampler before adopting its parameters for final VIX tables, and identify the sampler in optimizer checkpoints and table fingerprints. Defaults and previous mean-bridge results remain separately reproducible.

Two independent reciprocal states in smooth and rough time

This alternative assigns a separate reciprocal state to each stock-variance channel. The component subscripts in ζ_t and $\zeta_{\mathfrak{r}}$ identify their clocks, rather than time arguments. The affine driver and canonical boundary clock are

$$\begin{aligned}\mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)dt, & d\langle\xi\rangle &= 2\xi dt, \\ \mathfrak{r}_t &= \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^t \mathbf{1}_{\{\xi_s < \epsilon\}} ds, & 0 < \theta < 1.\end{aligned}\tag{4.145}$$

Here t is the deterministic internal time obtained from the configured calendar or business time map and its fitted expiry speeds. Market forwards, discounts and displayed implied volatilities retain their quoted calendar conventions. No inverse observation clock or affine stock-variance channel is added in this variant.

The complete discounted-stock specification is

$$\begin{aligned}\mathbb{E}_d[dS] &= 0, \\ \mathbb{E}_d[d\zeta_t] &= (\vartheta_t - \beta_t\zeta_t)dt, & d\langle\zeta_t\rangle &= 2\zeta_t dt, \\ \mathbb{E}_d[d\zeta_{\mathfrak{r}}] &= (\vartheta_{\mathfrak{r}} - \beta_{\mathfrak{r}}\zeta_{\mathfrak{r}})d\mathfrak{r}, & d\langle\zeta_{\mathfrak{r}}\rangle &= 2\zeta_{\mathfrak{r}} d\mathfrak{r}, \\ \frac{d\langle S \rangle}{S^2} &= \frac{\Sigma}{\zeta_t} dt + \frac{\Psi}{\zeta_{\mathfrak{r}}} d\mathfrak{r}, & & \\ \frac{d\langle\zeta_t, S\rangle}{S} &= \rho_t \sqrt{2\Sigma} dt, & \frac{d\langle\zeta_{\mathfrak{r}}, S\rangle}{S} &= \mathfrak{r} \sqrt{2\Psi} d\mathfrak{r}, \\ d\langle\zeta_t, \zeta_{\mathfrak{r}}\rangle &= d\langle\xi, \zeta_t\rangle = d\langle\xi, \zeta_{\mathfrak{r}}\rangle = d\langle\xi, S\rangle = 0.\end{aligned}\tag{4.146}$$

The rough state's quadratic variation uses $d\mathfrak{r}$, consistent with its purely rough dynamics. The displayed stock bracket is relative quadratic variation; the bracket of the level itself contains the factor S^2 . The two reciprocal Brownian drivers, their two residual stock Brownian drivers, and the affine clock driver are independent. Each stock component shares its reciprocal driver's innovation with correlation ρ_t or $\mathfrak{r}$. Thus these are reciprocal-state-stock correlations, positive for the requested negative variance-equity leverage. Both belong to $[0, 1]$ in this experiment; they are not the negative variance-stock correlations used in some earlier variants. The two reciprocal factors are independent, whereas ξ and $\zeta_{\mathfrak{r}}$ need not be independent: ξ determines the rough state's clock.

Numerator consistency and dimensions. With state diffusion fixed at two, the domestic constants ϑ_t and $\vartheta_{\mathfrak{r}}$ are themselves Feller characteristics. Stock

numeraire drift adds the stock–state bracket divided by S , so

$$\begin{aligned}\vartheta_{t,d} &= \vartheta_t, & \vartheta_{t,f} &= \vartheta_t + \rho_t \sqrt{2\Sigma}, \\ \vartheta_{\mathfrak{t},d} &= \vartheta_{\mathfrak{t}}, & \vartheta_{\mathfrak{t},f} &= \vartheta_{\mathfrak{t}} + \mathfrak{p} \sqrt{2\Psi}.\end{aligned}\tag{4.147}$$

The foreign characteristics are derived quantities, not additional free parameters. The inverse stock has the opposite component correlations in the stock measure. Reversion coefficients, state diffusions and the clock law are unchanged. Enforce $\zeta_{t,0}, \zeta_{\mathfrak{t},0} > 0$, $\Sigma, \Psi \geq 0$, $\kappa \geq 0$, $\xi_0 \geq 0$, and both domestic and foreign characteristics strictly greater than one. Equality at one prevents boundary attainment but gives divergent expected inverse-state variance, as discussed in the reciprocal-affine paper. No arbitrary extra Feller margin is imposed. Positive component correlations make the foreign inequalities automatic once the domestic inequalities hold.

If $[t] = T$, the normalized states and coefficients have dimensions

$$\begin{aligned}[\xi] &= [\zeta_t] = T, & [\mathfrak{t}] &= [\zeta_{\mathfrak{t}}] = T^{1-\theta}, \\ [\kappa] &= [\beta_t] = T^{-1}, & [\beta_{\mathfrak{t}}] &= T^{\theta-1}, \\ [\theta] &= [\vartheta_t] = [\vartheta_{\mathfrak{t}}] = [\rho_t] = [\mathfrak{p}] = [\Sigma] = [\Psi] = 1.\end{aligned}\tag{4.148}$$

In particular, both stock variance strengths are dimensionless. The normalization $d\langle\zeta\rangle = 2\zeta d(\text{own clock})$ fixes each reciprocal state’s scale; no separate reciprocal diffusion strength is fitted.

Simulation using the existing reciprocal step. First generate the ordinary affine rough-clock bank on the internal time grid using the established native-Torch sampled-touch and weighted bridge clock calculation. Independently evolve ζ_t over Δt and $\zeta_{\mathfrak{t}}$ over $\Delta \mathfrak{t}$. A zero rough-clock increment leaves the rough state and its stock contribution unchanged. No drift from the other clock is aggregated into either reciprocal state.

Reuse the existing exact noncentral Gamma endpoint and normalized conditional stock factor separately for each channel. The existing weighted reciprocal step receives state diffusion coefficient two, shape ϑ , reversion β , stock strength C and the negative of the public reciprocal-state correlation. It operates directly in the original state units. Only lanes with a positive own-clock increment enter the step; zero-increment lanes retain their exact current state and zero stock contribution. At $C = 0$, the existing exact state transition still operates, while its stock contribution is zero.

Smooth and rough factors use distinct seeded random streams. Each positive interval is split into exactly the configured number of equal own-clock increments: four for initial fitting and sixteen for final pricing. This fixed count is applied to every advancing path; no path-dependent financial subdivision cap is introduced.

The analytic endpoint multiplier has conditional mean one and matches the foreign first state moment. Its expected negative twice-log contribution is subtracted from the forward expected integrated reciprocal variance, leaving the established Gaussian bridge residual. Add the two conditional log-forwards and residual variances before evaluating option payoffs. Conditional independence preserves the analytically normalized martingale factor without sample centering. The endpoint laws are exact; the sampled rough clock and finite-step conditional-return/bridge approximation still require numerical convergence checks.

Initial calibration scope. The initial experiment uses the Monday 5 October 2026 end-of-day snapshot and the four SPXW expirations 6–9 October. Targets are reported closing premiums, with the established bid/ask-spread and volume weights, both rights and all supplied strikes. Each expiry interval has its own positive time speed; the first is fixed at one to anchor the overall time scale. The editable descriptor lists the thirteen model parameters explicitly, followed by the date-keyed speed dictionary. The separate implementation and notebook retain earlier model results. This specification does not assert a completed calibration or a fitted error.

Affine variance in the rough clock

This separate alternative retains the physical affine clock driver ξ and assigns a single affine variance state ζ to its boundary clock. The state ζ has no physical-time drift or diffusion. Both states contribute to physical stock variance, while ζ also contributes to rough stock variance. The domestic discounted-stock characteristics are

$$\begin{aligned}
\mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)dt, & d\langle\xi\rangle &= 2\xi dt, \\
\mathbb{E}_d[d\zeta] &= (\vartheta - \beta\zeta)d\mathfrak{r}, & d\langle\zeta\rangle &= 2\zeta d\mathfrak{r}, \\
\mathbb{E}_d[dS] &= 0, & d\langle\xi, \zeta\rangle &= 0, \\
\frac{d\langle S\rangle}{S^2} &= \frac{1}{2} [(\mathcal{H}\xi + \Sigma\zeta)dt + \boldsymbol{\Psi}\zeta d\mathfrak{r}], & & \\
\frac{d\langle\xi, S\rangle}{S} &= \rho\sqrt{\mathcal{H}}\xi dt, & \frac{d\langle\zeta, S\rangle}{S} &= \mathfrak{N}\sqrt{\boldsymbol{\Psi}}\zeta d\mathfrak{r}.
\end{aligned} \tag{4.149}$$

The factor $1/2$ applies to all three stock-variance weights. In particular, $\mathcal{H}$ retains the earlier Heston normalization and is distinct from a Hurst exponent. The canonical boundary clock is

$$\mathfrak{r}_t = \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^t \mathbf{1}_{\{\xi_s < \epsilon\}} ds, \quad 0 < \theta < 1. \tag{4.150}$$

Here t is the internal time obtained from the configured deterministic calendar or business map and any fitted expiry speeds. Market forwards, discounts and quoted implied-volatility maturities keep their original calendar conventions.

Covariances and attainable boundaries. Independent affine drivers and independent residual stock noises realize (4.149). The stock shares the ξ physical innovation with correlation ρ and the ζ rough innovation with correlation $\mathfrak{N}$, each in $[-1, 1]$. Since ζ is a variance state, negative $\mathfrak{N}$ gives negative equity leverage. This reverses the positive state–stock sign used for inverse variance in the preceding reciprocal alternative. The physical ξ –stock correlation, including the additional independent $\Sigma\zeta/2$ channel, is

$$\rho\sqrt{\frac{\mathcal{H}\xi}{\mathcal{H}\xi + \Sigma\zeta}}, \quad (4.151)$$

where the denominator is positive. There is no extra condition coupling ρ and $\mathfrak{N}$: they belong to different clock components. For the nonsticky affine boundary construction, dt and $d\mathfrak{t}$ are mutually singular. This explains the vanishing state cross-bracket and why the $\Sigma\zeta dt/2$ stock channel has no instantaneous ζ covariance. Dependence through the sampled rough clock still connects ξ and ζ .

The parameter domain is $\xi_0, \zeta_0, \vartheta, \kappa, \mathcal{H}, \Sigma, \boldsymbol{\psi} \geq 0$, with signed β and the correlation bounds above. No restriction $\vartheta \geq 1$ is imposed. For $0 < \vartheta < 1$ the rough affine state can touch zero; for $\vartheta = 0$ it is absorbed on reaching zero. The point mass at zero is part of its transition law. These cases do not create the inverse-state singularity of the reciprocal model. From a positive initial state, $\vartheta \geq 1$ prevents boundary attainment, but that is an optional regime, not a required stability constraint here.

Numeraire family and dimensions. Adding each stock–state covariance divided by S to the domestic drift gives the stock-measure parameters

$$\kappa_f = \kappa - \rho\sqrt{\mathcal{H}}, \quad \beta_f = \beta - \mathfrak{N}\sqrt{\boldsymbol{\psi}}, \quad \theta_f = \theta, \quad \vartheta_f = \vartheta. \quad (4.152)$$

These local drift identities leave the affine family intact. They also change the boundary-clock law through κ_f whenever the physical Heston leverage is active. The inverse stock has opposite channel correlations in the stock numeraire. The identities do not, by themselves, establish a global martingale theorem for every parameter choice.

Writing $[t] = T$, the fixed state diffusion coefficients imply

$$\begin{aligned} [\xi] &= T, & [\zeta] &= [\mathfrak{t}] = T^{1-\theta}, \\ [\kappa] &= T^{-1}, & [\beta] &= T^{\theta-1}, \\ [\mathcal{H}] &= T^{-2}, & [\Sigma] &= T^{\theta-2}, & [\boldsymbol{\psi}] &= T^{2\theta-2}. \end{aligned} \quad (4.153)$$

Shapes and correlations are dimensionless. The three stock coefficients are variance weights; unlike those of the preceding reciprocal model, they are not dimensionless. The six state and clock keys `theta`, `kappa`, `xi_0`, `zeta_0`, `vartheta`, and `beta` denote $\theta, \kappa, \xi_0, \zeta_0, \vartheta, \beta$, respectively. The five stock keys `H`, `Sigma`, `evan`, `rho`, and `ecorr` denote $\mathcal{H}, \Sigma, \Psi, \rho, \mathfrak{N}$.

Finite-grid simulation and stock normalization. Reuse the established native-Torch affine-driver and sampled boundary-clock bank. Apply the physical $\mathcal{H}$ stock step to the stored ξ endpoints of each bank interval. Split the interval into exactly N financial steps for ζ , using increments $\Delta t/N$ and $\Delta \mathfrak{t}/N$ on each path. The configured N is four for initial fitting and sixteen for final pricing; it is not a path-dependent cap. At zero rough increment the state is unchanged, although its independent physical stock variance remains. The affine endpoint transition is noncentral Gamma, including its Poisson–Gamma zero-shape atom. It uses native Torch draws on the existing ROCm backend.

For a generic own-clock affine endpoint with starting state x , shape a , reversion b and increment δ , write

$$\begin{aligned} c &= \frac{1 - e^{-b\delta}}{b}, & n &= xe^{-b\delta}, \\ \log M(q) &= -a \log(1 - cq) + \frac{nq}{1 - cq}, & 1 - cq &> 0, \end{aligned} \quad (4.154)$$

with $c = \delta$ at $b = 0$. The stock factor uses $\exp(qx_{\text{next}} - \log M(q))$, which has conditional mean one under this endpoint law. For stock variance $Cx/2$ in that clock and channel correlation r , the trapezoidal Itô coefficient is

$$q = \frac{r\sqrt{C}}{2} \left(1 + \frac{b\delta}{2} \right) - \frac{r^2 C \delta}{8}, \quad A = \frac{x + x_{\text{next}}}{2} \delta. \quad (4.155)$$

The independent residual variance is $(1 - r^2)CA/2$. The physical Heston inputs are $x = \xi$, $a = \theta$, $b = \kappa$, $C = \mathcal{H}$ and $r = \rho$. For the rough channel use $x = \zeta$, $a = \vartheta$, $b = \beta$, $C = \Psi$ and $r = \mathfrak{N}$. Add the independent physical variance $\Sigma A_{\zeta,t}/2$, where $A_{\zeta,t}$ sums the ζ trapezoids with physical increments. The full accumulated variance separately records all three weights with their factor $1/2$. The implementation's `fair_variance` field is this sampled trapezoidal integrated quadratic variation, not an evaluation of the conditional expectation derived below.

The inherited endpoint helper additionally requires $1 - 2cq > 0$ for a nondegenerate endpoint, so the multiplier has a finite second moment for ordinary Monte Carlo error estimates. This is a stated finite-step numerical condition, stronger than the first-moment domain in (4.154); an invalid trial is reported rather than

repaired by a sample normalization. The physical ξ normalizer uses its unconditional endpoint marginal, not a law conditioned on an internally sampled future touch. Conditional stock factors and residual lognormal factors then retain analytic normalization without empirical forward rescaling. Exact affine endpoint marginals and analytic normalization do not make the joint stock/clock transition exact. The endpoint-only tilt, sampled rough-clock law and trapezoidal areas remain finite-grid approximations. Increasing N alone does not refine the physical ξ bank grid.

Conditional fair variance. For an internal-time window Δt , the expected integrated relative quadratic variation is

$$\frac{1}{2}\mathbb{E}_t \left[\mathcal{H} \int_t^{t+\Delta t} \xi_s ds + \Sigma \int_t^{t+\Delta t} \zeta_s ds + \Psi \int_t^{t+\Delta t} \zeta_s d\mathfrak{r}_s \right]. \quad (4.156)$$

The ξ area is the ordinary affine mean area. Conditional on the future clock-driver path, the ζ mean is

$$\mathbb{E}[\zeta_{t+s} \mid \mathcal{F}_t, (\xi_{t+r})_{0 \leq r \leq s}] = \zeta_t e^{-\beta \Delta \mathfrak{r}_s} + \vartheta \frac{1 - e^{-\beta \Delta \mathfrak{r}_s}}{\beta}, \quad \Delta \mathfrak{r}_s = \mathfrak{r}_{t+s} - \mathfrak{r}_t. \quad (4.157)$$

Consequently the physical ζ area is obtained by integrating the clock average of this expression over $s \in [0, \Delta t]$. Its rough-clock area is

$$\mathbb{E}_t \left[\zeta_t \frac{1 - e^{-\beta \Delta \mathfrak{r}_{\Delta t}}}{\beta} + \frac{\vartheta}{\beta} \left(\Delta \mathfrak{r}_{\Delta t} - \frac{1 - e^{-\beta \Delta \mathfrak{r}_{\Delta t}}}{\beta} \right) \right]. \quad (4.158)$$

At $\beta = 0$, the conditional mean becomes $\zeta_t + \vartheta \Delta \mathfrak{r}_s$ and the rough area becomes $\zeta_t \mathbb{E}_t[\Delta \mathfrak{r}_{\Delta t}] + \vartheta \mathbb{E}_t[(\Delta \mathfrak{r}_{\Delta t})^2]/2$. Thus the fair-variance dependence on ζ_t is affine, but the future rough-clock average, including its dependence on ξ_t and κ , still has to be evaluated. Replacing the clock distribution by its mean is an approximation, not a consequence of affine state dynamics. This description asserts neither a calibrated option fit nor convergence of a VIX estimate; those require separate numerical results.

Affine variance in physical and rough time

The preceding affine-variance model extends by allowing the same variance state ζ to have affine drift and diffusion in physical time as well as in the boundary clock. The physical innovations of ξ and ζ have zero quadratic covariation. The

characteristics are

$$\begin{aligned}
\mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)dt, & d\langle\xi\rangle &= 2\xi dt, \\
\mathbb{E}_d[d\zeta] &= (\sigma_\zeta\vartheta_t - \beta_t\zeta)dt + (\vartheta - \beta\zeta)d\mathfrak{r}, & d\langle\zeta\rangle &= 2\zeta(\sigma_\zeta dt + d\mathfrak{r}), \\
\mathbb{E}_d[dS] &= 0, & d\langle\xi, \zeta\rangle &= 0, \\
\frac{d\langle S\rangle}{S^2} &= \frac{1}{2}[(\mathcal{H}\xi + \Sigma\zeta)dt + \boldsymbol{\Psi}\zeta d\mathfrak{r}], \\
\frac{d\langle\xi, S\rangle}{S} &= \rho\sqrt{\mathcal{H}}\xi dt, \\
\frac{d\langle\zeta, S\rangle}{S} &= \zeta\left[\rho_t\sqrt{\sigma_\zeta\Sigma}dt + \boldsymbol{\mathfrak{N}}\sqrt{\boldsymbol{\Psi}}d\mathfrak{r}\right].
\end{aligned} \tag{4.159}$$

The canonical clock is still (4.150), with $0 < \theta < 1$. The physical clock is the configured deterministic calendar or business-time map, including expiry speeds. The coefficient $\sigma_\zeta \geq 0$ scales the physical variance clock of ζ ; its physical shape is $\vartheta_t \geq 0$. If $\beta_t > 0$, physical dynamics alone revert toward $\sigma_\zeta\vartheta_t/\beta_t$. The existing coefficient Σ continues to weight physical stock variance. All three stock weights retain the factor $1/2$.

Correlation, boundaries and change of numeraire. Independent state innovations and independent residual stock innovations realize these covariances with $\rho, \rho_t, \boldsymbol{\mathfrak{N}} \in [-1, 1]$. The two total physical state–stock correlations, where defined, are

$$\rho\sqrt{\frac{\mathcal{H}\xi}{\mathcal{H}\xi + \Sigma\zeta}}, \quad \rho_t\sqrt{\frac{\Sigma\zeta}{\mathcal{H}\xi + \Sigma\zeta}}. \tag{4.160}$$

Their squared sum is at most one. Zero state covariation does not remove dependence through the boundary clock. Negative channel correlations give negative variance–stock leverage. Neither ϑ_t nor ϑ requires a lower bound of one; both are nonnegative and the variance state may reach zero. If both immigration terms vanish, zero is absorbing. Both reversion coefficients may have either sign. States and variance weights are nonnegative.

Adding the displayed stock covariances to the state drifts gives

$$\begin{aligned}
\kappa_f &= \kappa - \rho\sqrt{\mathcal{H}}, \\
\beta_{t,f} &= \beta_t - \rho_t\sqrt{\sigma_\zeta\Sigma}, \\
\beta_f &= \beta - \boldsymbol{\mathfrak{N}}\sqrt{\boldsymbol{\Psi}}.
\end{aligned} \tag{4.161}$$

Shapes and diffusion coefficients are unchanged. These are local stock-numeraire drift identities, not a global martingale theorem for arbitrary parameter choices.

Dimensions (4.153) remain valid, with the additions

$$[\sigma_\zeta] = T^{-\theta}, \quad [\beta_t] = T^{-1}, \quad [\vartheta_t] = [\rho_t] = 1. \quad (4.162)$$

The four new descriptor keys are

$$\begin{aligned} \sigma_\zeta & \text{ evar_zeta_physical} \\ \vartheta_t & \text{ vartheta_physical} \\ \beta_t & \text{ beta_physical} \\ \rho_t & \text{ rho_physical.} \end{aligned}$$

The preceding eleven keys retain their meanings. At $\sigma_\zeta = \beta_t = 0$ the preceding rough affine model is recovered, and the two other new parameters are inactive.

When the physical states combine. Suppress the rough-clock terms for this comparison and assume $\sigma_\zeta > 0$ and $\beta_t = \kappa$. Then

$$\begin{aligned} X &= \xi + \zeta/\sigma_\zeta, \\ \mathbb{E}_d[dX] &= (\theta + \vartheta_t - \kappa X)dt, \quad d\langle X \rangle = 2X dt. \end{aligned} \quad (4.163)$$

Thus the sum is a single normalized affine state. For the stock also to reduce to the same single Heston factor, require $\mathcal{H} = \sigma_\zeta \Sigma$ and $\rho_t = \rho$; then its variance and cross-covariance are $S^2 \mathcal{H} X dt/2$ and $S \rho \sqrt{\mathcal{H}} X dt$. Equal reversion alone are insufficient for this stock reduction. Setting only $\Psi = 0$ leaves rough evolution of ζ and does not produce this limit. The full model retains the boundary-clock dependence and two distinct states.

Affine substep. Retain the native-Torch path bank for the physical affine driver and its boundary clock. Apply the ξ stock step to its stored physical endpoints. Each bank interval is split into exactly the configured N financial substeps, without path-dependent step counts. For each substep define

$$\begin{aligned} D &= \sigma_\zeta \Delta t + \Delta \mathfrak{r}, & A &= \sigma_\zeta \vartheta_t \Delta t + \vartheta \Delta \mathfrak{r}, \\ B &= \beta_t \Delta t + \beta \Delta \mathfrak{r}, & V &= \Sigma \Delta t + \Psi \Delta \mathfrak{r}, \\ K &= \rho_t \sqrt{\sigma_\zeta \Sigma} \Delta t + \mathfrak{n} \sqrt{\Psi} \Delta \mathfrak{r}. \end{aligned} \quad (4.164)$$

These are integrated coefficients, not additional model parameters. For $D > 0$, use the affine endpoint law over increment D with shape A/D , reversion B/D and stock variance coefficient V/D . Its effective correlation is $K/\sqrt{DV}$ when $V > 0$, and zero at $V = 0$. Cauchy–Schwarz gives $K^2 \leq DV$, including the independent residual stock noises, so this is an admissible channel correlation. Use (4.154) and (4.155) with these inputs; accumulate the full stock variance separately from the

correlation-reduced residual variance. At $D = 0$, ζ evolves by its deterministic physical decay; there is no stochastic state endpoint to sample. Its endpoint and physical area are $\zeta e^{-\beta_t \Delta t}$ and $\zeta(1 - e^{-\beta_t \Delta t})/\beta_t$, respectively, with continuous area limit $\zeta \Delta t$ at $\beta_t = 0$.

The conditional endpoint factor is normalized analytically. The inherited check on its second moment is retained. Invalid finite-step trials are reported without empirical centering, sample rescaling or added financial bounds. The zero-extension case uses the preceding simulator unchanged. The endpoint law is exact for the aggregated substep, not for the original joint random-clock dynamics. In general physical and rough affine generators do not commute. Aggregation, sampled clock increments, the endpoint-only stock tilt and state-area quadrature require grid-convergence checks. The configured initial and final financial substep counts are four and sixteen. Increasing them does not refine the physical ξ bank grid.

Conditional variance structure. The fair-variance decomposition (4.156) still applies. Conditional on the future ξ path, define for $s \geq 0$

$$\begin{aligned} C_s &= \beta_t s + \beta(\mathfrak{r}_{t+s} - \mathfrak{r}_t), \\ m_s &= e^{-C_s} \left[\zeta_t + \int_0^s e^{C_r} (\sigma_\zeta \vartheta_t dr + \vartheta d\mathfrak{r}_{t+r}) \right]. \end{aligned} \quad (4.165)$$

Here m_s is the conditional ζ mean. The expected physical and rough areas are respectively the clock-path expectations of $\int_0^{\Delta t} m_s ds$ and $\int_0^{\Delta t} m_s d\mathfrak{r}_{t+s}$. Both remain affine in current ζ_t , but depend on the future clock path, not generally just its terminal value or its mean. No calibrated VIX accuracy or variance convergence follows without numerical checks.

Bilinear feedback through the rough affine variance state

This extension returns to the rough-only ζ dynamics of Section 4.1. The additional physical drift and diffusion of ζ in Section 4.1 are absent. Instead, ζ changes the physical Heston variance weight and the mean reversion of the clock driver. The two added coefficients are $\mathcal{H}_\zeta$ and κ_ζ , corresponding to the code names `H_zeta` and `kappa_zeta`. The domestic discounted-stock characteristics are

$$\begin{aligned} \mathbb{E}_d[d\xi] &= [\theta - (\kappa + \kappa_\zeta \zeta)\xi] dt, & d\langle \xi \rangle &= 2\xi dt, \\ \mathbb{E}_d[d\zeta] &= (\vartheta - \beta\zeta)d\mathfrak{r}, & d\langle \zeta \rangle &= 2\zeta d\mathfrak{r}, \\ \mathbb{E}_d[dS] &= 0, & d\langle \xi, \zeta \rangle &= 0, \\ \frac{d\langle S \rangle}{S^2} &= \frac{1}{2} [((\mathcal{H} + \mathcal{H}_\zeta \zeta)\xi + \Sigma\zeta)dt + \boldsymbol{\Psi}\zeta d\mathfrak{r}], & & \\ \frac{d\langle \xi, S \rangle}{S} &= \rho\sqrt{\mathcal{H} + \mathcal{H}_\zeta \zeta} \xi dt, & \frac{d\langle \zeta, S \rangle}{S} &= \boldsymbol{\mathfrak{n}}\sqrt{\boldsymbol{\Psi}} \zeta d\mathfrak{r}. \end{aligned} \quad (4.166)$$

The factor $1/2$ is retained for every stock-variance contribution, including $\mathcal{H}_\zeta \xi \zeta dt$. The new drift contribution is $-\kappa_\zeta \xi \zeta dt$. These are separate parameters; neither coefficient determines the other. Setting both to zero gives (4.149).

The boundary clock retains its canonical normalization,

$$\mathfrak{r}_t = \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^t \mathbf{1}_{\{\xi_s < \epsilon\}} ds, \quad 0 < \theta < 1, \quad (4.167)$$

where ξ now has the feedback drift in (4.166). The configured deterministic business clock and expiry-speed map still determine t from calendar time. Market forwards, discounts and implied-volatility maturities retain their source conventions. There is no magic clock in this specification.

Correlation convention and boundaries. The physical stock channel with variance $(\mathcal{H} + \mathcal{H}_\zeta \zeta) \xi dt/2$ shares the physical ξ innovation with correlation ρ . Its residual noise and the $\Sigma \zeta dt/2$ stock noise are independent of both state drivers. The rough stock channel shares the rough ζ innovation with correlation $\mathfrak{N}$. Thus the total physical ξ –stock correlation, where its denominator is positive, is

$$\rho \sqrt{\frac{(\mathcal{H} + \mathcal{H}_\zeta \zeta) \xi}{(\mathcal{H} + \mathcal{H}_\zeta \zeta) \xi + \Sigma \zeta}}. \quad (4.168)$$

The state cross-bracket remains zero despite the dependence created by feedback. For the nonsticky boundary construction, dt and $d\mathfrak{r}$ are mutually singular, so the two channel correlations need no additional constraint coupling them.

Nonnegative variance over the whole unbounded state domain requires $\mathcal{H}, \mathcal{H}_\zeta, \Sigma, \vartheta \geq 0$. Initial states and ϑ are nonnegative, with $|\rho|, |\mathfrak{N}| \leq 1$. The reuse of the existing nonnegative-reversion clock sampler requires $\kappa, \kappa_\zeta \geq 0$; this is an identified sampler domain, not a claim that negative reversion cannot define an SDE. The rough reversion β remains signed. There is no reciprocal variance singularity and no restriction $\vartheta \geq 1$: $\vartheta = 0$ includes the absorbing zero-state atom. These local conditions alone are not a proof of global nonexplosion or stock martingality for every feedback parameter choice. In particular, stronger mean reversion of ξ also changes the speed of the clock that drives ζ .

Stock-measure drift and dimensions. Adding the stock–state covariance divided by S to each domestic drift gives the local stock-measure characteristics

$$\begin{aligned} \mathbb{E}_f[d\xi] &= \left[\theta - \left(\kappa + \kappa_\zeta \zeta - \rho \sqrt{\mathcal{H} + \mathcal{H}_\zeta \zeta} \right) \xi \right] dt, \\ \mathbb{E}_f[d\zeta] &= \left[\vartheta - (\beta - \mathfrak{N} \sqrt{\vartheta}) \zeta \right] d\mathfrak{r}. \end{aligned} \quad (4.169)$$

The constant drifts and state quadratic variations are unchanged. Unless the square-root term degenerates, the first line is not a reversion coefficient linear in ζ . Consequently this extension is not closed under the earlier constant affine-parameter numeraire map. No additional drift term is introduced to force such closure.

The original dimensions in (4.153) remain in force. The additional ones follow by matching the existing variance and reversion weights:

$$[\mathcal{H}_\zeta] = T^{\theta-3}, \quad [\kappa_\zeta] = T^{\theta-2}, \quad [\mathcal{H}_\zeta \zeta] = T^{-2}, \quad [\kappa_\zeta \zeta] = T^{-1}. \quad (4.170)$$

The model parameter list is the original eleven entries together with `H_zeta` and `kappa_zeta`. It excludes `evan_zeta_physical`, `varthetaeta_physical`, `beta_physical`, and `rho_physical` from the preceding smooth- ζ experiment.

Accepted numerical plan and carried first contact. An independent precomputed ξ and boundary-clock bank is no longer valid when $\kappa_\zeta \neq 0$: future clock increments depend on the evolving ζ state. The accepted treatment advances the existing physical intervals sequentially. At each interval start it freezes ζ in the two coefficients $\kappa + \kappa_\zeta \zeta$ and $\mathcal{H} + \mathcal{H}_\zeta \zeta$. The user's correction is to retain the first contact in the mean-reversion-adjusted coordinate, rather than retaining an absolute physical contact time or resampling an unconsumed contact.

Define

$$K_t = \int_0^t (\kappa + \kappa_\zeta \zeta_s) ds, \quad q_t = \int_0^t e^{K_s} ds, \quad Y_{q_t} = e^{K_t} \xi_t. \quad (4.171)$$

The time-changed state satisfies $dY = \theta dq + \sqrt{2Y} dB_q$. Its first-contact target q_{hit} is retained until reached. At a physical interval start, the equivalent local residual is

$$r_t = e^{-K_t} (q_{\text{hit}} - q_t). \quad (4.172)$$

For the frozen rate $k = \kappa + \kappa_\zeta \zeta_t$ and interval δ , the local transformed duration and untouched residual update are

$$a = \frac{e^{k\delta} - 1}{k}, \quad r_{t+\delta} = e^{-k\delta} (r_t - a) \quad \text{if } r_t > a, \quad (4.173)$$

with $a = \delta$ at $k = 0$. A contact occurs within the interval when $r_t \leq a$; its physical offset is $\log(1 + kr_t)/k$, interpreted as r_t at $k = 0$. The local residual representation avoids forming large exponentials of accumulated K_t . A newly sampled contact has local residual ξ_t/G , $G \sim \Gamma(1 - \theta, 1)$. Keeping a target for accumulated K_t alone is not equivalent to keeping q_{hit} when the rate changes: q integrates e^K .

The native-Torch contact-conditioned affine endpoints and the existing weighted boundary-clock approximation are to be used with this carried target. The configured number of rough- ζ financial steps is unchanged. Existing internal clock

substeps retain their numerical condition $k\delta \leq 1/2$, now evaluated for each path's frozen rate; this is not an added financial parameter bound. The observation grid, backend, quote target and endpoint-stock formulas are retained.

This frozen-coefficient joint step is an approximation of the coupled dynamics, not an exact joint transition. Its drift, variance and conditional stock normalizer must all use the same frozen coefficients. The carried contact is latent future information, not an observable conditioning variable for stock pricing. The stock endpoint tilt therefore retains the unconditional affine endpoint cumulant given the observable starting state, not the cumulant of the contact-conditioned bridge. Forcing conditional stock mean one given the future contact could remove the intended stock–contact dependence. Implementation must verify the joint transition, contact retention, observable-state stock mean and raw forward errors; a mean-one assertion in the enlarged contact filtration would not suffice. The implementation advances the coupled states jointly and therefore does not reuse an independent clock bank across ζ -parameter trials. Tests cover retained contacts, endpoint moments, observable-state stock means and exact recovery of the preceding model at zero feedback. These checks do not establish convergence of the finite-grid coupled approximation; calibration and raw forward diagnostics are reported separately.

Fair variance. The conditional expected integrated relative stock variance over $[t, t + \Delta t]$ is

$$\frac{1}{2} \mathbb{E}_t \left[\int_t^{t+\Delta t} ((\mathcal{H} + \mathcal{H}_\zeta \zeta_s) \xi_s + \Sigma \zeta_s) ds + \boldsymbol{\Psi} \int_t^{t+\Delta t} \zeta_s d\mathbf{r}_s \right]. \quad (4.174)$$

The feedback makes the future clock depend on the variance-state innovations. Therefore the earlier conditional- ζ formulas based on a clock generated independently of those innovations cannot simply be reused. This expectation requires a calculation for the coupled process. No fair-variance accuracy or calibration result is asserted here.

A power of rough affine variance in the physical stock channel

This extension of Section 4.1 replaces only the physical stock-variance contribution $\Sigma \zeta dt/2$ by $\Sigma \zeta^\gamma dt/2$, with $\gamma > 0$. The state dynamics, clock, correlations and rough stock-variance contribution are unchanged. In particular, ζ still evolves entirely in the boundary clock; there is no added physical ζ drift or diffusion and no $\xi \zeta$ feedback in the clock driver. The domestic discounted-stock characteristics

are

$$\begin{aligned}
\mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)dt, & d\langle\xi\rangle &= 2\xi dt, \\
\mathbb{E}_d[d\zeta] &= (\vartheta - \beta\zeta)d\mathfrak{r}, & d\langle\zeta\rangle &= 2\zeta d\mathfrak{r}, \\
\mathbb{E}_d[dS] &= 0, & d\langle\xi, \zeta\rangle &= 0, \\
\frac{d\langle S\rangle}{S^2} &= \frac{1}{2}[(\mathcal{H}\xi + \Sigma\zeta^\gamma)dt + \boldsymbol{\Psi}\zeta d\mathfrak{r}], \\
\frac{d\langle\xi, S\rangle}{S} &= \rho\sqrt{\mathcal{H}}\xi dt, & \frac{d\langle\zeta, S\rangle}{S} &= \mathfrak{n}\sqrt{\boldsymbol{\Psi}}\zeta d\mathfrak{r}.
\end{aligned} \tag{4.175}$$

All stock-variance weights retain the factor $1/2$. The rough clock is

$$\mathfrak{r}_t = \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^t \mathbf{1}_{\{\xi_s < \epsilon\}} ds, \quad 0 < \theta < 1. \tag{4.176}$$

Internal physical time t retains the configured deterministic calendar/business map and expiry-speed settings. Source forwards, discounts and implied-volatility maturities keep their market conventions.

Correlations, measures and the zero boundary. The $\Sigma\zeta^\gamma$ stock channel is independent of both state innovations, just as the original $\Sigma\zeta$ channel. Therefore no new state-stock covariance or numeraire drift is introduced. The total physical ξ -stock correlation, where defined, is now

$$\rho \sqrt{\frac{\mathcal{H}\xi}{\mathcal{H}\xi + \Sigma\zeta^\gamma}}. \tag{4.177}$$

The stock-measure drift identities remain

$$\kappa_f = \kappa - \rho\sqrt{\mathcal{H}}, \quad \beta_f = \beta - \mathfrak{n}\sqrt{\boldsymbol{\Psi}}, \quad \theta_f = \theta, \quad \vartheta_f = \vartheta. \tag{4.178}$$

These are local measure-change identities, not a global martingale theorem for every parameter regime.

The base parameter domain is unchanged, with the sole additional condition $\gamma > 0$. In particular, $\vartheta \geq 0$ has no lower Feller restriction of one, and β is signed. The value at zero is exactly $0^\gamma = 0$; there is no state floor. Fractional powers can have a singular derivative at zero, so ordinary smooth Itô formulas for the transformed state must not be used across an attainable boundary without justification. Pricing does not transform the ζ SDE: it evaluates the nonnegative power only in the stock-variance integral. The exact absorbing atom at $\vartheta = 0$ is therefore preserved.

Dimensions and a level-matched comparison. The new code parameter `gamma` denotes the dimensionless exponent γ . The original eleven parameter names and their meanings are retained. Since $[\zeta] = T^{1-\theta}$ and stock variance has units T^{-1} , only the dimension of Σ changes:

$$[\gamma] = 1, \quad [\Sigma] = T^{-1-\gamma(1-\theta)}. \quad (4.179)$$

At $\gamma = 1$ this reduces to $T^{\theta-2}$ and the original model is recovered exactly. At a positive reference state, a comparison that holds the physical stock-variance contribution fixed may start with

$$\Sigma_{\text{new}} = \Sigma_{\text{old}} \zeta_0^{1-\gamma}, \quad \Sigma_{\text{new}} \zeta_0^\gamma = \Sigma_{\text{old}} \zeta_0. \quad (4.180)$$

This is a transparent parameter initialization, not a sample-dependent price adjustment or a constraint during fitting. Relative to that matched level, $0 < \gamma < 1$ raises the contribution at smaller positive ζ and compresses it at larger ζ . It does not remove an atom at zero. Over a narrow observed state range, changes in Σ and γ can partly compensate, so a small objective change alone does not establish that the exponent has been identified.

Finite-grid calculation. Keep the base model’s native-Torch affine transitions, ROCm backend, physical grid, sampled boundary-clock bank and fixed number N of financial steps. At a ζ substep with physical length δt and endpoints ζ_i, ζ_{i+1} , replace only the physical independent stock variance by

$$\Delta V_\Sigma = \frac{\Sigma \delta t}{4} (\zeta_i^\gamma + \zeta_{i+1}^\gamma). \quad (4.181)$$

This is the trapezoid of the endpoint powers, not the power of a state trapezoid. The correlated ξ and rough- ζ stock factors retain their existing analytic affine normalizers. For the independent physical channel, conditional on the state path and its variance V_Σ , the stock factor $\exp(-V_\Sigma/2 + \sqrt{V_\Sigma}Z)$ has mean one for an independent standard normal Z . Conditional lognormal payoff integration therefore remains unchanged. No empirical normalization or new sampler is required. At $\gamma = 1$ the original implementation is the exact numerical nesting target, including random draws and prices.

The clock approximation, endpoint stock tilt and variance quadrature remain finite-grid approximations. Initial fitting keeps the established lower-path procedure and four financial steps; confirmed raw pricing uses at least 2^{20} paths, sixteen financial steps and no control variate. Every fitted expiration retains its forward and parity checks. This specification does not assert that implementation tests or the new calibration have completed.

Conditional fair variance. The conditional expected integrated relative stock variance is

$$\frac{1}{2}\mathbb{E}_t\left[\mathcal{H}\int_t^{t+\Delta t}\xi_s ds + \Sigma\int_t^{t+\Delta t}\zeta_s^\gamma ds + \mathfrak{V}\int_t^{t+\Delta t}\zeta_s d\mathfrak{r}_s\right]. \quad (4.182)$$

Only its physical- ζ component changes. For a given future clock increment the affine ζ transition is unchanged, but the required quantity is its fractional moment $\mathbb{E}[\zeta_s^\gamma \mid \mathcal{F}_t, \Delta\mathfrak{r}_s]$, not $(\mathbb{E}[\zeta_s \mid \mathcal{F}_t, \Delta\mathfrak{r}_s])^\gamma$. For $0 < \gamma < 1$, concavity makes the latter an upper bound when the expectations exist. The result must then also be averaged over the future clock distribution. The physical contribution is generally no longer affine in the initial ζ . No additional fair-variance approximation or VIX convergence claim is introduced here.

Physical affine variance with logarithmic rough drift

The physical affine driver ξ generates the canonical boundary clock $\mathfrak{r}$. A second state ζ has physical square-root diffusion and power diffusion in that rough clock. The discounted stock has one physical variance contribution from each state and a rough variance contribution from ζ . Its domestic characteristics are

$$\begin{aligned} \mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)dt, & d\langle\xi\rangle &= 2\xi dt, \\ \mathbb{E}_d[d\zeta] &= (\vartheta - \kappa_\zeta\zeta)dt - \mathfrak{I}\zeta^\gamma \log(\zeta/\mathfrak{V})d\mathfrak{r}, \\ d\langle\zeta\rangle &= 2\zeta dt + \mathfrak{V}\zeta^{2\gamma-1}d\mathfrak{r}, \\ \mathbb{E}_d[dS] &= 0, & d\langle\xi, \zeta\rangle &= 0, \\ \frac{d\langle S\rangle}{S^2} &= (H_\xi\xi + H_\zeta\zeta)dt + \mathfrak{I}\zeta d\mathfrak{r}, \\ \frac{d\langle\xi, S\rangle}{S} &= \rho_\xi\sqrt{2H_\xi}\xi dt, \\ \frac{d\langle\zeta, S\rangle}{S} &= \rho_\zeta\sqrt{2H_\zeta}\zeta dt + \mathfrak{M}\sqrt{\mathfrak{I}\mathfrak{V}}\zeta^\gamma d\mathfrak{r}. \end{aligned} \quad (4.183)$$

In particular, the physical second stock term is $H_\zeta\zeta$, not a constant, and its covariance is linear in ζ . There is no overall factor $1/2$ in the displayed stock variance. The clock normalization is

$$\mathfrak{r}_t = \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^t \mathbf{1}_{\{\xi_s < \epsilon\}} ds, \quad 0 < \theta < 1. \quad (4.184)$$

The correlations refer to the corresponding stock-variance channels. The physical state noises are orthogonal. Taking the channel correlations in $[-1, 1]$ gives a positive semidefinite covariance matrix; a channel correlation need not equal the correlation with the total stock noise.

Change of numeraire. The local stock-numeraire drift adds each state–stock cross-bracket in (4.183). For $\mathfrak{J} > 0$, this stays in the same drift family under the parameter changes

$$\begin{aligned}\kappa_{\mathfrak{f}} &= \kappa - \rho_{\xi} \sqrt{2H_{\xi}}, \\ \kappa_{\zeta, \mathfrak{f}} &= \kappa_{\zeta} - \rho_{\zeta} \sqrt{2H_{\zeta}}, \\ \mathfrak{y}_{\mathfrak{f}} &= \mathfrak{y} \exp\left(\frac{\mathfrak{m} \sqrt{\mathfrak{n} \psi}}{\mathfrak{J}}\right).\end{aligned}\tag{4.185}$$

All shapes, diffusion coefficients and powers are unchanged. When $\mathfrak{J} = 0$, the additional rough drift is $\mathfrak{m} \sqrt{\mathfrak{n} \psi} \zeta^{\gamma} d\mathfrak{t}$ and cannot generally be represented by a finite target shift while keeping $\mathfrak{J} = 0$. These are local drift identities. True martingality of the continuous stock additionally requires the relevant integrability and boundary conditions; it does not follow solely from writing a locally normalized numerical stock factor.

Parameters and boundary qualifications. Require $\gamma, \zeta_0, \mathfrak{y} > 0$ and non-negative $\kappa, \xi_0, \vartheta, H_{\xi}, H_{\zeta}, \mathfrak{n}, \psi, \mathfrak{J}$; κ_{ζ} may be signed. No additional Feller floor is imposed on ϑ . These sign restrictions are not a general existence theorem for every power. The physical square-root component can reach zero when its Feller characteristic is below one; logarithmic drift acting only in $\mathfrak{r}$ does not automatically prohibit such a physical-time hit. Although $\zeta^{\gamma} \log(\zeta/\mathfrak{y})$ tends to zero for $\gamma > 0$, the separately frozen coefficients below can be undefined at zero. A numerical candidate with an undefined active coefficient is rejected, not repaired by a positive state floor. This is a restriction of that approximation at the encountered state, not an asserted financial bound. Inactive-clock and zero-coefficient terms are omitted before evaluation.

The physical normalization $d\langle \zeta \rangle = 2\zeta dt + \dots$ implies $[\xi] = [\zeta] = T$, while $[\mathfrak{r}] = T^{1-\theta}$. Thus

$$\begin{aligned}[H_{\xi}] &= [H_{\zeta}] = T^{-2}, & [\kappa] &= [\kappa_{\zeta}] = T^{-1}, \\ [\mathfrak{n}] &= T^{\theta-2}, & [\psi] &= T^{\theta+2-2\gamma}, \\ [\mathfrak{J}] &= T^{\theta-\gamma}, & [\mathfrak{y}] &= T.\end{aligned}\tag{4.186}$$

The shapes, correlations and exponent γ are dimensionless, and the logarithm has the dimensionless argument $\zeta/\mathfrak{y}$. The descriptor names are `theta`, `kappa`, `xi_0`, `zeta_0`, `vartheta`, `kappa_zeta`, `H_xi`, `H_zeta`, `he`, `ev`, `gamma`, `nun`, `ayin`, `rho_xi`, `rho_zeta` and `ecorr`, respectively.

The local affine approximation. The implementation keeps the existing physical affine-driver/clock bank, native-Torch endpoint sampler and configured

physical grid. The financial dynamics update once per physical interval, using its full physical and rough-clock increments. At the interval's starting state ζ_i , freeze

$$q_i = \frac{\Psi}{2} \zeta_i^{2\gamma-2}, \quad b_i = \mathfrak{I} \zeta_i^{\gamma-1} \log(\zeta_i/\mathfrak{Y}). \quad (4.187)$$

With the full interval increments δt and $\delta \mathfrak{t}$, form

$$\begin{aligned} D_i &= \delta t + q_i \delta \mathfrak{t}, & A_i &= \vartheta \delta t, \\ B_i &= \kappa_\zeta \delta t + b_i \delta \mathfrak{t}, & V_i &= 2(H_\zeta \delta t + \mathfrak{I} \delta \mathfrak{t}), \\ K_i &= \rho_\zeta \sqrt{2H_\zeta} \delta t + \mathfrak{D} \sqrt{2\mathfrak{I} q_i} \delta \mathfrak{t}. \end{aligned} \quad (4.188)$$

For $D_i > 0$, use the existing affine transition of duration D_i , shape A_i/D_i and rate B_i/D_i . This reproduces the requested infinitesimal coefficients at ζ_i . The transition is exact for the frozen affine step, not for the full nonlinear dynamics. The stock tilt has effective correlation $K_i/\sqrt{D_i V_i}$ when defined. Cauchy-Schwarz gives $K_i^2 \leq D_i V_i$ under the channel correlation bounds. The factor two in V_i converts the displayed stock variance to the established helper's factor-1/2 convention.

The stock endpoint multiplier is chosen to match the affine endpoint mean under the stock numeraire. Consider the local affine transition

$$\begin{aligned} dX &= (a - bX)du + \sqrt{2X} dW, \\ \frac{d\langle S \rangle}{S^2} &= wX du, & \frac{d\langle X, S \rangle}{S} &= r\sqrt{2w} X du. \end{aligned} \quad (4.189)$$

Its stock-measure rate is $b_f = b - r\sqrt{2w}$. Write $\text{decay}(x) = (1 - e^{-x})/x$, continued at zero by $\text{decay}(0) = 1$. For a duration δu , set

$$\begin{aligned} c &= \delta u \text{decay}(b\delta u), & n &= X_0 e^{-b\delta u}, \\ m_d &= n + ac, & m_f &= X_0 e^{-b_f \delta u} + a\delta u \text{decay}(b_f \delta u). \end{aligned} \quad (4.190)$$

The domestic endpoint log moment-generating function and the stock multiplier are

$$\begin{aligned} K(q) &= -a \log(1 - cq) + \frac{nq}{1 - cq}, \\ q &= \frac{2(m_f - m_d)}{c \left(2m_f - ac + \sqrt{(ac)^2 + 4nm_f} \right)}, \\ R(X_1) &= \exp(qX_1 - K(q)). \end{aligned} \quad (4.191)$$

For the nondegenerate affine transition, this chooses the moment-generating domain $cq < 1$ and solves $K'(q) = m_f$. Consequently

$$\mathbb{E}_d[R] = 1, \quad \mathbb{E}_d[X_1 R] = m_f, \quad (4.192)$$

with initial state and local coefficients fixed. Zero-duration, inactive-stock and absorbing-zero cases retain the limit $R = 1$. The endpoint factor is an approximation to the conditional stock return matching these two moments, not its exact conditional law.

For the ξ contribution, use $(a, b, w, r) = (\theta, \kappa, H_\xi, \rho_\xi)$ and $\delta u = \delta t$. For the combined ζ contribution, use $(a, b, w, r) = (A_i/D_i, B_i/D_i, V_i/(2D_i), K_i/\sqrt{D_i V_i})$ and $\delta u = D_i$. Thus the combined stock-measure rate is $(B_i - K_i)/D_i$, or integrated rate $B_i - K_i$. The moment identities for ξ use its affine marginal endpoint distribution; they do not assert the same identities conditional on the rough-clock increment generated jointly with that endpoint. For ζ , they apply to its frozen affine transition with the sampled interval inputs fixed. They do not establish exact finite-step stock moments for the full nonlinear model.

The conditional residual lognormal variance is unchanged and its factor has mean one. Neither factor uses a sample mean or an empirical forward correction. Trapezoidal integration is retained for each interval's physical and rough fair-variance contributions, separately:

$$\delta V_{\zeta,t} = H_\zeta \delta t \frac{\zeta_i + \zeta_{i+1}}{2}, \quad \delta V_{\zeta,\mathfrak{r}} = \mathfrak{n} \delta \mathfrak{r} \frac{\zeta_i + \zeta_{i+1}}{2}. \quad (4.193)$$

Degenerate inactive channels retain the deterministic limits. The full increments are used in one update; this physical-plus-rough-clock model has no magic-time integration or financial-substep parameter. Clock sampling, freezing, endpoint tilt and area integration remain finite-grid approximations whose convergence must be checked on the configured grid.

Implementation correction, 2026-10-16. A four/sixteen subdivision setting had been carried over from magic-time models without authorization. That setting is removed for this model. Earlier calibration and pricing comparisons using it are retained as historical evidence, not as results for the corrected interval update; those parameters and comparisons require repricing. The subsequent endpoint-factor correction replaces the normalized trapezoidal exponent by (4.191), matching both moments in (4.192) for each local affine law. It does not change the residual variance or the configured grid. Prices and timings recorded before this endpoint correction likewise require repricing.

Calibration and variance diagnostics. The separate notebook displays these dynamics before its editable parameter descriptor. A nested expiration-date dictionary retains the expiry-speed parameters: the first speed is fixed at one, and the remaining speeds are available to the joint fit. This changes internal dynamics without changing the quoted forwards, discounts or option implied-volatility

maturity convention. The initial comparison uses the 2026-10-05 reported closing prices and all 33 expirations through 2026-12-31, within the explicitly configured three-month limit. The established volume/bid/ask target includes calls and puts, ITM and OTM, at every supplied strike with positive weight.

Fitting uses direct stock simulation with Torch sorting and cumulative sums to price all supplied strikes. Confirmation and all displayed fit plots instead use conditional payoff integration with 2^{20} paths. The current and retained training errors, confirmed error and evaluation counts are distinguished. A provisional direct-pricing improvement does not replace the confirmed plot set; the latter remains visible while a candidate is checked.

The fitting grid starts with four quarter-day intervals, two half-day intervals and two one-day intervals, followed by intervals of two, three, four and five days, repeating five-day intervals thereafter. These durations are measured in business-adjusted physical time before applying the fitted expiration-speed map, and every expiration is inserted as an exact grid point. The exposed fitting stages use $2^{15}, 2^{17}, 2^{20}$ paths with maximum intervals of five, two and one adjusted days, respectively. The confirmation grid has intervals no longer than one adjusted day on that same pre-speed axis. Each resulting interval receives one state update; there is no inner magic-time subdivision. A configurable trigger, initially two unsuccessful confirmations at the current fitting precision, permits early promotion to the next finer fitting stage.

Every result states its estimator, path count, grid and control policy. The existing forward control variate may be used in both training and confirmation. Controlled and raw normalized errors remain separate, and the shared ROCm GPU runs one calculation at a time. For every expiry, raw forward errors with their Monte Carlo standard errors, and raw parity errors, are computed from the unadjusted simulation without controls. These diagnostics and the simulated stock paths are not changed by applying the payoff control.

Conditional fair variance would require the expectation of

$$\int_t^{t+\Delta t} (H_\xi \xi_s + H_\zeta \zeta_s) ds + \mathfrak{n} \int_t^{t+\Delta t} \zeta_s d\mathfrak{s}_s. \quad (4.194)$$

The logarithmic drift and power diffusion do not imply an affine conditional expectation for ζ . No closed-form fair-variance approximation, VIX result or new calibration improvement is asserted by this specification.

Symmetric logarithmic variance

This section records the model specification and numerical procedure. The combined conditional-Gamma construction below is implemented and calibration is

ongoing; post-fit grid convergence remains to be checked. The preceding construction distinguishes the clock driver from the variance state. A symmetric extension gives both states physical square-root diffusion and a boundary clock, and lets each clock act on the other state. The indices are fixed as follows: λ_i belongs to a power of ξ_i , and each rough coefficient has the index of the clock multiplying it. For the discounted stock in the domestic measure,

$$\begin{aligned}\mathbb{E}_d[d\xi_1] &= (\theta_1 - \kappa_1 \xi_2^{\lambda_2} \xi_1)dt - \mathfrak{I}_2 \xi_1^{\gamma_2} \log(\xi_1/\mathfrak{Y}_2) d\mathfrak{R}_2, \\ \mathbb{E}_d[d\xi_2] &= (\theta_2 - \kappa_2 \xi_1^{\lambda_1} \xi_2)dt - \mathfrak{I}_1 \xi_2^{\gamma_1} \log(\xi_2/\mathfrak{Y}_1) d\mathfrak{R}_1, \\ d\langle \xi_1 \rangle &= 2\xi_1 dt + \mathfrak{V}_2 \xi_1^{2\gamma_2-1} d\mathfrak{R}_2, \\ d\langle \xi_2 \rangle &= 2\xi_2 dt + \mathfrak{V}_1 \xi_2^{2\gamma_1-1} d\mathfrak{R}_1, \\ d\langle \xi_1, \xi_2 \rangle &= 0, \quad \mathbb{E}_d[dS] = 0.\end{aligned}\tag{4.195}$$

The stock variance and state–stock covariances are

$$\begin{aligned}\frac{d\langle S \rangle}{S^2} &= \left(H_1 \xi_2^{2\lambda_2} \xi_1 + H_2 \xi_1^{2\lambda_1} \xi_2 \right) dt + \mathfrak{I}_2 \xi_1 d\mathfrak{R}_2 + \mathfrak{I}_1 \xi_2 d\mathfrak{R}_1, \\ \frac{d\langle \xi_1, S \rangle}{S} &= \rho_1 \sqrt{2H_1} \xi_2^{\lambda_2} \xi_1 dt + \mathfrak{I}_2 \sqrt{\mathfrak{I}_2 \mathfrak{V}_2} \xi_1^{\gamma_2} d\mathfrak{R}_2, \\ \frac{d\langle \xi_2, S \rangle}{S} &= \rho_2 \sqrt{2H_2} \xi_1^{\lambda_1} \xi_2 dt + \mathfrak{I}_1 \sqrt{\mathfrak{I}_1 \mathfrak{V}_1} \xi_2^{\gamma_1} d\mathfrak{R}_1.\end{aligned}\tag{4.196}$$

Thus ρ_i and $\mathfrak{I}_j$ are correlations with their respective stock channels. Independent state noises and channel correlations in $[-1, 1]$ produce a positive semidefinite instantaneous covariance. A channel correlation is not the correlation with the total stock noise. There is no overall factor one half in the stock variance.

Clocks and the corner. The intended normalization of each boundary clock is

$$\mathfrak{r}_i(t) = \lim_{\epsilon \downarrow 0} \frac{\theta_i(1 - \theta_i)}{\epsilon^{\theta_i}} \int_0^t \mathbf{1}_{\{\xi_i(s) < \epsilon\}} ds, \quad \mathbf{1}_{\{\xi_i > 0\}} d\mathfrak{R}_i = 0.\tag{4.197}$$

The proposed domain is

$$0 < \theta_1 < 1, \quad 0 < \theta_2 < 1, \quad \theta_1 + \theta_2 > 1.\tag{4.198}$$

Each state is intended to reach its own boundary, while simultaneous zeros are intended to be excluded. For this mutually rough system the last inequality is an imposed condition, not an established corner classification: the independent affine-pair argument does not by itself cover the additional boundary drift and diffusion. Existence of the coupled clocks and their joint boundary behavior require verification.

Local numeraire closure. The factors $2\lambda_j$ in stock variance make the physical state–stock drift adjustment proportional to the existing reversion term. For $i \neq j$ and $\mathfrak{J}_j > 0$, the stock-numeraire changes are

$$\theta_i^S = \theta_i, \quad \kappa_i^S = \kappa_i - \rho_i \sqrt{2H_i}, \quad \mathfrak{Y}_j^S = \mathfrak{Y}_j \exp\left(\frac{\mathfrak{N}_j \sqrt{\mathfrak{N}_j \mathfrak{P}_j}}{\mathfrak{J}_j}\right). \quad (4.199)$$

The remaining state and variance coefficients are unchanged. Quoting $1/S$ under the stock numeraire reverses the stock-channel correlations. These are local drift identities, not a proof of a valid global numeraire change in every parameter regime. If $\mathfrak{J}_j = 0$ and the rough covariance is nonzero, the added $\xi_i^{\gamma_j} d\mathfrak{r}_j$ drift cannot be absorbed into a finite shifted $\mathfrak{Y}_j$; closure at that special case requires a separate coefficient of that power drift.

Dimensions. With $[\xi_i] = T$ and $[\mathfrak{r}_i] = T^{1-\theta_i}$, the dimensions are

$$\begin{aligned} [\theta_i] = [\lambda_i] = [\gamma_i] = [\rho_i] = [\mathfrak{N}_i] &= 1, \\ [\kappa_i] &= T^{-1-\lambda_j}, \quad [H_i] = T^{-2-2\lambda_j}, \\ [\mathfrak{Y}_j] &= T, \quad [\mathfrak{J}_j] = T^{\theta_j-\gamma_j}, \\ [\mathfrak{N}_j] &= T^{\theta_j-2}, \quad [\mathfrak{P}_j] = T^{\theta_j+2-2\gamma_j}, \quad i \neq j. \end{aligned} \quad (4.200)$$

The target in the logarithm therefore has the same dimension as the state on which its clock acts. Changing a shape or a power changes the dimensions of the corresponding dimensional coefficients.

Signed feedback powers. The exponents λ_1, λ_2 may be negative. For $\lambda_i < 0$, the corresponding initial state $\xi_{i,0}$ must be strictly positive. The existing frozen-coefficient transition is retained; no state floor, coefficient cap, or replacement boundary rule is added. Negative powers are singular at zero, so permitting them numerically does not establish global existence or boundary integrability of the coupled model. In particular, $\theta_i + 2\lambda_i > 0$ is a local affine occupation-integrability heuristic, not a proved admissibility theorem for this system and not an imposed parameter bound.

Requested fresh-touch approximation. At each configured outer interval, start from the actual current pair and freeze $\kappa_i \xi_j^{\lambda_j}$ in the local physical affine touch calculation. Draw fresh candidate touch times, with feedback from a touch of the other state inactive in that local calculation. Use the temporary draws only to determine the interval’s rough-clock increments, complete one coupled state update, and discard the candidate times. The next interval redraws from its updated

pair; it does not retain an independent clock bank or a countdown from a previous interval.

Combined conditional-Gamma transition. Let C_i denote this interval's physical contact record, and retain the existing conditional noncentral-Gamma parameters (c_i, s_i, r_i) : shape, scale, and noncentral mean. Their moment-generating function is

$$G_i(v \mid C_i) = (1 - s_i v)^{-c_i} \exp\left(\frac{r_i v}{1 - s_i v}\right). \quad (4.201)$$

The ordinary physical CIR endpoint MGF is $G_i(v) = \mathbb{E}[G_i(v \mid C_i)]$. Let u_i be the existing physical stock tilt that matches the local stock-measure endpoint mean. Freeze the nonlinear rough coefficients at the current state pair and combine the full interval inputs:

$$\begin{aligned} q_i &= \tfrac{1}{2} \mathfrak{V}_j \xi_i^{2\gamma_j-2}, & b_i &= \mathfrak{I}_j \xi_i^{\gamma_j-1} \log(\xi_i / \mathfrak{Y}_j), \\ D_i &= s_i + q_i \delta \mathfrak{T}_j, & A_i &= c_i s_i, \\ B_i &= b_i \delta \mathfrak{T}_j, & J_i &= \mathfrak{N}_j \sqrt{2\mathfrak{N}_j q_i} \delta \mathfrak{T}_j, & i &\neq j. \end{aligned} \quad (4.202)$$

Define $\phi(B) = (1 - e^{-B})/B$, with $\phi(0) = 1$. Draw one endpoint X_i from the noncentral-Gamma law with shape A_i/D_i , scale $D_i\phi(B_i)$, and noncentral mean $r_i e^{-B_i}$; use the deterministic limits where necessary. Its log MGF is

$$K_i(v) = -\frac{A_i}{D_i} \log(1 - D_i\phi(B_i)v) + \frac{r_i e^{-B_i} v}{1 - D_i\phi(B_i)v}. \quad (4.203)$$

Physical conditioning is retained in (c_i, s_i, r_i) rather than replaced by an unconditional CIR endpoint. The combined draw is not a physical update followed by a rough update.

The physical conditional Esscher inputs and the local stock-measure target are

$$r_i^S = \frac{r_i}{(1 - s_i u_i)^2}, \quad A_i^S = \frac{c_i s_i}{1 - s_i u_i}, \quad m_i^S = r_i^S e^{-(B_i - J_i)} + A_i^S \phi(B_i - J_i). \quad (4.204)$$

Choose v_i satisfying $K_i'(v_i) = m_i^S$ using the existing closed-form Gamma mean-matching solve. Retain the physical likelihood factor carried by the contact record:

$$W_i = \frac{G_i(u_i \mid C_i)}{G_i(u_i)}, \quad R_i = \exp(v_i X_i - K_i(v_i)), \quad R = W_1 W_2 R_1 R_2. \quad (4.205)$$

Conditional on the interval-start pair, the physical contact records are independent, and the endpoint Gamma draws are independent given both records. Therefore

$$\begin{aligned} \mathbb{E}[W_i] &= 1, & \mathbb{E}[R_i \mid C_1, C_2] &= 1, & \mathbb{E}[R] &= 1, \\ \mathbb{E}[X_i R \mid C_1, C_2] &= W_1 W_2 m_i^S. \end{aligned} \quad (4.206)$$

These identities are analytic properties of the local transition, not normalizations estimated from the simulated sample. At $\delta\mathfrak{f}_j = 0$, the merged Gamma parameters reduce exactly to (c_i, s_i, r_i) , $v_i = u_i$, and

$$W_i R_i = \frac{\exp(u_i X_i)}{G_i(u_i)}. \quad (4.207)$$

Thus the original physical endpoint factor is recovered exactly. No new quadrature is required for this construction; the existing native Torch Gamma and clock primitives are retained.

There is one combined state update per configured outer interval, no magic clock, and no internal subdivision. The frozen coefficients and independently proposed contact records remain a finite-step approximation to mutual feedback. In particular, m_i^S is the local approximation's foreign-mean target, not an assertion of the exact finite-step mean of the continuous nonlinear SDE. Distinct touches inside one interval are retained and are not interpreted as simultaneous zero states. Undefined active coefficients must be reported rather than replaced by a state floor or silently clipped.

Direct stock sampling at requested expirations. The direct estimator accumulates the conditional log-forward L and residual variance V along the configured state intervals. It draws one independent stock residual normal per requested expiration rather than per outer interval. For successive requested expirations T_k ,

$$M_k = M_{k-1} + \sqrt{V_k - V_{k-1}} Z_k, \quad \log(S_{T_k}/S_0) = L_k - \frac{1}{2}V_k + M_k, \quad (4.208)$$

where $M_0 = V_0 = 0$ and the Z_k are independent standard normals, independent of the state and clock draws. Carrying M_k forward preserves the conditional residual dependence across maturities. Independently redrawing the entire accumulated residual at each expiry would not do so. This requested sampling convention leaves the state paths and grid unchanged and is implemented with the combined transition above.

Initial calibration experiment. The separate three-month experiment uses 2026-10-05 end-of-day reported closes and all 33 available expirations through 2026-12-31, within the explicit 2027-01-05 maturity limit. The established volume/bid/ask target includes positive-weight calls and puts, ITM and OTM, at every supplied strike. A nested expiration-date speed dictionary has one fixed unit anchor. Internal time rescaling does not change quoted forwards, discounts, or implied-volatility maturities.

The intended existing fitting procedure uses direct stock simulation and native Torch sorting and cumulative sums. Progressive fitting samples use 2^{15} , 2^{17} , and

2^{20} paths. Ongoing confirmation and published fit plots use conditional payoff integration with the explicitly configured 2^{19} paths. The user permits smaller samples when their Monte Carlo errors are sufficiently small; the path count must remain explicit. Final pricing and the later time-step convergence study retain at least 2^{20} paths. The configured grid begins with quarter-day outer intervals and becomes coarser later; confirmation intervals are at most one day on the configured pre-speed axis. There are no internal substeps. One calibration at a time uses the shared AMD/ROCm GPU through Jupyter. The starting experiment retains the existing identity calendar-time base clock, followed by the fitted expiration-speed map.

The existing payoff control variate is allowed, while raw forward errors, their Monte Carlo standard errors, and raw parity diagnostics are computed from the unadjusted simulation. Current and retained training errors, confirmation error, evaluation counts, and the actual parameter descriptor remain visible. Every fitted expiration has an implied-volatility and volume plot. The new model's code, checkpoints and results are separate from its asymmetric predecessor; no prior calibration score is transferred without repricing these dynamics.

Time-step convergence after a reasonable fit. After obtaining a reasonable fit, hold all model parameters and the fitted expiry speeds fixed. Use at least 2^{20} paths and the same seed, progressively refine the outer intervals, and compare option prices, the calibration target error, raw forward errors with their Monte Carlo standard errors, and both clocks' distributions and touch frequencies at each maturity. A common seed does not imply identical state paths on different grids. This later check is performed without recalibration; it does not change the current fitting settings or introduce internal subdivisions.

Symmetric reciprocal rough drift

This separate variant preserves the two states ξ_1, ξ_2 and their boundary clocks, while replacing the logarithmic rough drift by a reciprocal-affine drift. The first experiment fixes $\lambda_1 = \lambda_2 = 0$ and $\gamma_1 = \gamma_2 = 2$. The new coefficients $\mathfrak{J}_i$ and ϑ_i are indexed by the state on which they act. The diffusion and stock-channel coefficients $\mathfrak{V}_j, \mathfrak{N}_j, \mathfrak{M}_j$ remain indexed by the opposite clock, where $j = 3 - i$. Under the domestic measure,

$$\begin{aligned}\mathbb{E}_d[d\xi_i] &= (\theta_i - \kappa_i \xi_i)dt + \left[\mathfrak{J}_i \xi_i + \mathfrak{V}_j \left(1 - \frac{\vartheta_i}{2} \right) \xi_i^2 \right] d\mathfrak{R}_j, \\ d\langle \xi_i \rangle &= 2\xi_i dt + \mathfrak{V}_j \xi_i^3 d\mathfrak{R}_j, \quad i \neq j, \\ d\langle \xi_1, \xi_2 \rangle &= 0, \quad \mathbb{E}_d[dS] = 0.\end{aligned}\tag{4.209}$$

The stock coefficients, with the same channel-correlation convention, are

$$\begin{aligned}\frac{d\langle S \rangle}{S^2} &= (H_1\xi_1 + H_2\xi_2)dt + \mathfrak{n}_2\xi_1 d\mathfrak{r}_2 + \mathfrak{n}_1\xi_2 d\mathfrak{r}_1, \\ \frac{d\langle \xi_i, S \rangle}{S} &= \rho_i\sqrt{2H_i}\xi_i dt + \mathfrak{n}_j\sqrt{\mathfrak{n}_j\psi_j}\xi_i^2 d\mathfrak{r}_j, \quad i \neq j.\end{aligned}\tag{4.210}$$

The two physical state channels and two rough state channels are orthogonal; $|\rho_i|, |\mathfrak{n}_j| \leq 1$ gives a positive semidefinite instantaneous covariance. These correlations refer to individual stock channels, not to the aggregate stock innovation. The boundary-clock normalization remains (4.197). The imposed domain $0 < \theta_i < 1$, $\theta_1 + \theta_2 > 1$ has the same intended corner interpretation and the same unresolved coupled-boundary qualification as (4.198).

Reciprocal state and Feller characteristics. For the rough component alone, Itô's formula gives

$$\begin{aligned}\mathbb{E}_d[d(1/\xi_i)]|_{d\mathfrak{r}_j} &= \left(\frac{\psi_j\vartheta_i}{2} - \frac{\mathfrak{z}_i}{\xi_i}\right) d\mathfrak{r}_j, \\ d\langle 1/\xi_i \rangle|_{d\mathfrak{r}_j} &= \frac{\psi_j}{\xi_i} d\mathfrak{r}_j.\end{aligned}\tag{4.211}$$

Thus ϑ_i is the Feller characteristic of the affine reciprocal. Equivalently, $2/(\psi_j\xi_i)$ has constant drift ϑ_i and quadratic variation twice its level in clock $\mathfrak{r}_j$, when $\psi_j > 0$. Positive $\mathfrak{z}_i$ means positive linear rough drift of ξ_i and mean reversion of its reciprocal. The coefficient $\mathfrak{z}_i$ may also be negative. This reciprocal is affine only in the isolated rough component, not in physical time.

The local stock-numeraire parameter change is

$$\begin{aligned}\theta_i^S &= \theta_i, \quad \kappa_i^S = \kappa_i - \rho_i\sqrt{2H_i}, \\ \mathfrak{z}_i^S &= \mathfrak{z}_i, \quad \vartheta_i^S = \vartheta_i - 2\mathfrak{n}_j\sqrt{\mathfrak{n}_j/\psi_j}, \quad i \neq j.\end{aligned}\tag{4.212}$$

The formula for ϑ_i^S applies to active rough diffusion, $\psi_j > 0$. The affine reciprocal avoids zero when its Feller characteristic is at least one; a finite inverse first moment requires it to be strictly greater than one. The numerical model therefore requires

$$\vartheta_i > 1, \quad \vartheta_i^S > 1 \quad \text{when } \psi_j > 0,\tag{4.213}$$

without an additional arbitrary margin. These are rough-component conditions, not a global existence or martingale theorem for the coupled model. If $\psi_j = 0$, the associated rough covariance and quadratic drift vanish, ϑ_i is inactive, and no division by ψ_j is used. Inverting the stock quote under its own numeraire reverses the channel correlations.

Dimensions and starting-point map. With $[\xi_i] = T$ and $[\mathfrak{r}_j] = T^{1-\theta_j}$,

$$\begin{aligned} [\kappa_i] &= T^{-1}, & [H_i] &= T^{-2}, \\ [\mathfrak{J}_i] &= T^{\theta_j-1}, & [\mathfrak{P}_j] = [\mathfrak{r}_j] &= T^{\theta_j-2}. \end{aligned} \quad (4.214)$$

The shapes and correlations are dimensionless. Unlike the preceding model, the new $\mathfrak{J}_i$ is state indexed and has a different time dimension. There is no logarithmic target $\mathfrak{Y}_j$ in this variant. For a positive starting state and active rough diffusion, the previous logarithmic fit gives the initialization

$$\mathfrak{J}_i^{\text{new}} = \mathfrak{J}_j^{\text{old}} \xi_{i,0}, \quad \vartheta_i = 2 + \frac{2\mathfrak{J}_j^{\text{old}}}{\mathfrak{P}_j} \left[1 + \log \left(\frac{\xi_{i,0}}{\mathfrak{Y}_j^{\text{old}}} \right) \right]. \quad (4.215)$$

It matches the old rough drift and its first state derivative at $\xi_{i,0}$ after setting the old power to two. It does not assert equivalence of the two full models, particularly when the old fitted γ_i and λ_i differ from two and zero. A zero old logarithmic drift maps to $\mathfrak{J}_i = 0$, $\vartheta_i = 2$. Inadmissible mapped parameters are reported rather than clipped.

Retained numerical procedure. The implementation uses the combined conditional-Gamma construction (4.201)– (4.207) with only the following frozen rough coefficients changed:

$$q_i = \tfrac{1}{2} \mathfrak{P}_j \xi_i^2, \quad b_i = -\mathfrak{J}_i + \mathfrak{P}_j \left(\frac{\vartheta_i}{2} - 1 \right) \xi_i. \quad (4.216)$$

Fresh physical contact records retain each state’s conditional endpoint law. Each full opposite-clock increment enters the same combined transition once. The analytic contact likelihood and mean-matched endpoint stock factor retain their mean-one identity for that local law. This remains a frozen-coefficient approximation; the affine form of the reciprocal rough block does not make the combined physical/rough transition exact. No reciprocal endpoint split, new quadrature, internal subdivision, or path batching is added. Direct residual stock sampling remains (4.208).

Separate calibration comparison. The first comparison retains the 2026-10-05 close target and its 33 expirations through 2026-12-31, within the explicit three-month limit 2027-01-05. Only $\vartheta_1, \vartheta_2, \mathfrak{J}_1, \mathfrak{J}_2$ are fitted initially. Every other coefficient and all retained expiry speeds are fixed, with the explicit exceptions that $\lambda_i = 0$ and $\gamma_i = 2$ define the requested new variant. The existing volume/bid/ask target, all positive-weight strikes and both option rights, Torch estimators, outer

grid, and raw-forward diagnostics are preserved. Ongoing confirmations use the configured 2^{19} paths; final pricing and subsequent fixed-parameter outer-grid convergence use at least 2^{20} . The implementation, notebook, checkpoints, and results are separate from the logarithmic model. The old model's retained score is an initialization reference only; this section asserts neither a completed new fit nor grid convergence.

Symmetric power coordinates for smooth and rough variance

There are two states, each expressed in a smooth coordinate ξ_i and a rough coordinate ζ_i . For $i = 1, 2$ and $j = 3 - i$, set

$$\zeta_i = \mathfrak{D}_i \xi_i^{\lambda_i}, \quad \mathfrak{D}_i > 0, \quad \lambda_i \neq 0. \quad (4.217)$$

Both signs of λ_i are allowed. These coordinate exponents are unrelated to the preceding smooth-feedback exponents. In particular, the old choice of zero cannot be retained here: it would make the coordinate transform constant. The four symbols represent two states, not four independent diffusions.

Clock components and covariances. The defining affine variance normalizations apply to different clock components. The rough mean-reversion rate may depend on the state:

$$\begin{aligned} \mathbb{E}_d[d\xi_i]|_{dt} &= (\theta_i - \kappa_i \xi_i)dt, & d\langle \xi_i \rangle|_{dt} &= 2\xi_i dt, \\ \mathbb{E}_d[d\zeta_i]|_{d\mathfrak{r}_j} &= (\vartheta_i - \mathfrak{D}_i \xi_i^{\gamma_i} \zeta_i)d\mathfrak{r}_j, & d\langle \zeta_i \rangle|_{d\mathfrak{r}_j} &= 2\zeta_i d\mathfrak{r}_j. \end{aligned} \quad (4.218)$$

Here $\mathfrak{r}_j$ is the boundary clock of the other smooth coordinate, with normalization (4.197). The new exponents $\gamma_i \in \mathbb{R}$ multiply the rough mean-reversion rate. The multiplier is ξ_i , the state other than the clock driver ξ_j , which is zero on the support of $d\mathfrak{r}_j$. Setting $\gamma_1 = \gamma_2 = 0$ recovers the preceding two-factor SDE. The corrected numerical composition below need not reproduce the preceding merged finite-step approximation. These exponents are distinct from the coordinate powers λ_i and from the former logarithmic-model exponents. The domain $0 < \theta_i < 1$, $\theta_1 + \theta_2 > 1$ retains its intended exclusion of simultaneous zeros; as above, sufficiency for the fully coupled boundary-clock system is not asserted. There is no magic clock.

Under the domestic measure the discounted stock has zero drift and

$$\begin{aligned}\frac{d\langle S \rangle}{S^2} &= \frac{H_1 \xi_1 + H_2 \xi_2}{2} dt + \frac{\mathfrak{V}_1}{2\zeta_1} d\mathfrak{r}_2 + \frac{\mathfrak{V}_2}{2\zeta_2} d\mathfrak{r}_1, \\ \frac{d\langle \xi_i, S \rangle}{S} \Big|_{dt} &= \rho_i \sqrt{H_i} \xi_i dt, \\ \frac{d\langle \zeta_i, S \rangle}{S} \Big|_{d\mathfrak{r}_j} &= \mathfrak{N}_i \sqrt{\mathfrak{V}_i} d\mathfrak{r}_j, \quad d\langle \xi_1, \xi_2 \rangle = 0.\end{aligned}\tag{4.219}$$

All coefficients in this variant are state indexed. In particular, $\mathfrak{V}_i$ scales the rough stock channel associated with ζ_i , and $\mathfrak{N}_i$ is that coordinate's stock-channel correlation. The correlation of the corresponding ξ_i rough noise is instead $\text{sign}(\lambda_i)\mathfrak{N}_i$. With independent state noises and $|\rho_i|, |\mathfrak{N}_i| \leq 1$, these are component Brownian correlations and give positive semidefinite covariance. Positive $\mathfrak{N}_i$ corresponds to negative rough variance–stock leverage.

The full ξ equations on the positive interior follow from Itô's formula:

$$\begin{aligned}\mathbb{E}_d[d\xi_i] &= (\theta_i - \kappa_i \xi_i) dt + \left[\frac{\vartheta_i - 1 + 1/\lambda_i}{\lambda_i \mathfrak{D}_i} \xi_i^{1-\lambda_i} - \frac{\mathfrak{J}_i}{\lambda_i} \xi_i^{1+\gamma_i} \right] d\mathfrak{r}_j, \\ d\langle \xi_i \rangle &= 2\xi_i dt + \frac{2}{\lambda_i^2 \mathfrak{D}_i} \xi_i^{2-\lambda_i} d\mathfrak{r}_j, \\ \frac{d\langle \xi_i, S \rangle}{S} &= \rho_i \sqrt{H_i} \xi_i dt + \frac{\mathfrak{N}_i \sqrt{\mathfrak{V}_i}}{\lambda_i \mathfrak{D}_i} \xi_i^{1-\lambda_i} d\mathfrak{r}_j.\end{aligned}\tag{4.220}$$

Conversely, the induced smooth terms of ζ_i are

$$\begin{aligned}\mathbb{E}_d[d\zeta_i]_{dt} &= \left[\lambda_i(\theta_i + \lambda_i - 1) \frac{\zeta_i}{\xi_i} - \lambda_i \kappa_i \zeta_i \right] dt, \\ d\langle \zeta_i \rangle_{dt} &= 2\lambda_i^2 \frac{\zeta_i^2}{\xi_i} dt, \\ \frac{d\langle \zeta_i, S \rangle}{S} &= \lambda_i \rho_i \sqrt{H_i} \zeta_i dt + \mathfrak{N}_i \sqrt{\mathfrak{V}_i} d\mathfrak{r}_j.\end{aligned}\tag{4.221}$$

Consequently neither coordinate is generally affine in the full combined dynamics. For negative λ_i , the rough stock coefficient $1/\zeta_i$ vanishes as ξ_i tends to zero; for positive λ_i it diverges. The intended boundary protection in the latter case is inactivity of the opposite clock. This support observation alone does not establish integrability near the corner.

Numeraire changes and dimensions. The parameter family is locally closed under the stock-numeraire drift change:

$$\kappa_i^S = \kappa_i - \rho_i \sqrt{H_i}, \quad \vartheta_i^S = \vartheta_i + \mathfrak{N}_i \sqrt{\mathfrak{V}_i}.\tag{4.222}$$

The parameters $\theta_i, \lambda_i, \gamma_i, \mathfrak{D}_i, \mathfrak{I}_i$ are unchanged. An inverse stock quote reverses the channel correlations. At $\gamma_i = 0$, the pure rough affine channel has the usual zero-boundary threshold one, and finite inverse first moment requires its Feller characteristic to exceed one. The comparison retains $\vartheta_i > 1$ and $\vartheta_i^S > 1$, with no arbitrary additional margin. These rough-channel requirements and the local drift identity do not prove global existence or martingality for the coupled model. For nonzero γ_i , the additional rough-coordinate drift is

$$-\mathfrak{I}_i \mathfrak{D}_i^{-\gamma_i/\lambda_i} \zeta_i^{1+\gamma_i/\lambda_i}. \quad (4.223)$$

It vanishes at $\zeta_i = 0$ when $1 + \gamma_i/\lambda_i > 0$. At exponent zero, it shifts the constant boundary drift to $\vartheta_i - \mathfrak{I}_i \mathfrak{D}_i$, with the corresponding stock-measure shift as well. At a negative exponent it is singular unless $\mathfrak{I}_i = 0$. Consequently the retained conditions $\vartheta_i, \vartheta_i^S > 1$ alone are not sufficient in those latter regimes. Both signs of γ_i are permitted without an arbitrary exponent bound; a finite numerical step does not establish boundary stability or finiteness of moments.

The normalization fixes the dimensional assignments

$$\begin{aligned} [\xi_i] &= T, & [\mathfrak{I}_j] &= [\zeta_i] = T^{1-\theta_j}, \\ [\kappa_i] &= T^{-1}, & [H_i] &= T^{-2}, \\ [\mathfrak{D}_i] &= T^{1-\theta_j-\lambda_i}, & [\mathfrak{I}_i] &= T^{\theta_j-1-\gamma_i}, & [\boldsymbol{\Psi}_i] &= 1. \end{aligned} \quad (4.224)$$

The shape parameters, exponents, correlations and expiry speeds are dimensionless. In particular, the new rough stock coefficient $\boldsymbol{\Psi}_i$ has a different meaning and dimension from the preceding clock-indexed rough state-variance coefficient.

Exact nesting of the preceding reciprocal coefficients. For positive old rough state-variance coefficients, the preceding reciprocal variant is recovered at $\lambda_i = -1, \gamma_i = 0$ by

$$\mathfrak{D}_i = \frac{2}{\boldsymbol{\Psi}_j^{\text{old}}}, \quad \boldsymbol{\Psi}_i = \frac{4\mathfrak{I}_j^{\text{old}}}{\boldsymbol{\Psi}_j^{\text{old}}}, \quad \mathfrak{I}_i = -\mathfrak{I}_j^{\text{old}}, \quad H_i = 2H_i^{\text{old}}, \quad j = 3 - i. \quad (4.225)$$

All initial states, smooth shapes and rates, rough shapes and rates, smooth correlations and expiry speeds remain unchanged. The factor two in H_i accounts for the one-half stock-variance convention in (4.219). This is an exact SDE and frozen coefficient identity; it does not imply equality of prices computed with different finite-step schemes or a completed new calibration.

Four sign regimes. The comparison treats $(\text{sign } \lambda_1, \text{sign } \lambda_2) \in \{(-, -), (-, +), (+, -), (+, +)\}$ as separate calibration regimes. Zero is a singular coordinate map, so a local fit is

not expected to connect these regimes continuously. For a coordinate sign reversal at a positive starting state, keeping γ_i fixed, a matched initial point is obtained from

$$\lambda'_i = -\lambda_i, \quad \mathfrak{D}'_i = \mathfrak{D}_i \xi_{i,0}^{2\lambda_i}, \quad \mathfrak{J}'_i = -\mathfrak{J}_i + \frac{2(\vartheta_i - 1)}{\zeta_{i,0} \xi_{i,0}^{\gamma_i}}. \quad (4.226)$$

All other coefficients remain unchanged. This preserves $\zeta_{i,0}$, the local state drift and variance, and the rough stock variance at the initial state. At $\gamma_i = 0$ it is exactly the previous sign-flip map. It does not preserve the whole process or option prices. Keeping $\mathfrak{D}_i$ as the ζ_i -stock correlation reverses the rough ξ_i -stock covariance when the power sign changes. Each regime has its own sign-preserving exponent descriptors, checkpoints, calibration notebook and live page; there is no finite bound on exponent magnitude. Each page reports that regime's own current and retained errors, parameters and plots, without substituting another sign regime's result. Calibrations run sequentially on the GPU.

Smooth-coordinate and reciprocal-coordinate simulation. The numerical revision `conditional_physical_xi_then_equilibrium_rough_zeta_v2` uses smooth transitions in ξ and rough transitions in ζ . For each configured outer interval, first redraw the two contact records with rates κ_i and draw both smooth endpoints ξ_i^{sm} using the existing touch-conditioned affine law. The established smooth conditional stock factor, analytic contact likelihood and residual variance retain the stock weight $H_i/2$. After both smooth draws, the rough block starts from

$$\zeta_i^{\text{sm}} = \mathfrak{D}_i (\xi_i^{\text{sm}})^{\lambda_i}, \quad u = \Delta \mathfrak{r}_j, \quad j = 3 - i. \quad (4.227)$$

A zero opposite-clock increment is the identity and requires no coordinate powers. For positive u , the treatment of the rough drift depends on the existing coefficients, without adding a fitted parameter.

When $q_i = \gamma_i/\lambda_i > 0$ and $\mathfrak{J}_i > 0$, set

$$q_i = \frac{\gamma_i}{\lambda_i}, \quad a_i = \mathfrak{J}_i \mathfrak{D}_i^{-q_i}, \quad \zeta_{*,i} = \left(\frac{\vartheta_i}{a_i} \right)^{1/(1+q_i)}, \quad K_i = \frac{\vartheta_i}{\zeta_{*,i}}. \quad (4.228)$$

The rough drift has the exact decomposition

$$\vartheta_i - a_i \zeta_i^{1+q_i} = (\vartheta_i - K_i \zeta_i) + (K_i \zeta_i - a_i \zeta_i^{1+q_i}). \quad (4.229)$$

Both deterministic pieces have zero drift at $\zeta_{*,i}$. This level is not the stationary mean of the stochastic process. The second piece admits the exact flow

$$F_{i,h}(z) = \left[e^{-q_i K_i h} z^{-q_i} + \frac{a_i}{K_i} (1 - e^{-q_i K_i h}) \right]^{-1/q_i}. \quad (4.230)$$

The rough update is

$$\widehat{\zeta}_i = F_{i,u/2}(\zeta_i^{\text{sm}}), \quad Y_i \sim \text{CIR}(\widehat{\zeta}_i, u; \vartheta_i, K_i, 2), \quad \zeta_i^{\text{next}} = F_{i,u/2}(Y_i), \quad (4.231)$$

where the CIR notation denotes the retained affine transition with drift $\vartheta_i - K_i Y$ and quadratic variation $2Y du$. The deterministic maps are evaluated in logarithmic form. They surround one existing noncentral Gamma draw; they do not add outer intervals, rough-clock samples or random transitions. As q_i tends to zero from above, K_i tends to $\mathfrak{I}_i$ and $F_{i,h}$ to the identity. Thus this correction approaches the original exact affine rough step continuously.

For $\gamma_i = 0$, the original exact affine rough step with rate $\mathfrak{I}_i$ is retained directly. If $\mathfrak{I}_i = 0$, its zero-rate transition is retained. All other cases outside the corrected branch retain the previous rate

$$b_i = \mathfrak{I}_i(\xi_i^{\text{sm}})^{\gamma_i}$$

frozen over the rough interval. In particular, the correction does not resolve the accuracy or boundary behaviour of $q_i < 0$ or negative $\mathfrak{I}_i$. A damping-only split at $q_i < 0$ could artificially reach zero even when the complete rough drift repels the boundary; it is not substituted into those regimes.

For notation common to both branches, put $b_i = K_i$ in (4.231). In the unchanged branch put $\widehat{\zeta}_i = \zeta_i^{\text{sm}}$ and $\zeta_i^{\text{next}} = Y_i$. The CIR endpoint has shape ϑ_i , scale $u\phi(b_i u)$ and noncentral mean $\widehat{\zeta}_i e^{-b_i u}$, where $\phi(v) = (1 - e^{-v})/v$ and $\phi(0) = 1$. In either branch the state is then mapped back by

$$\xi_i^{\text{next}} = \left(\frac{\zeta_i^{\text{next}}}{\mathfrak{I}_i} \right)^{1/\lambda_i}.$$

Both states use the same smooth-then-rough staging, and contact records are redrawn from the resulting states at the next configured outer interval. There are no extra physical-time subdivisions or magic-time substeps.

The existing reciprocal conditional stock-mean approximation supplies a rough stock factor R_i from the intermediate CIR endpoint Y_i , not from its final nonlinear image. Its analytic normalization and CIR stock-measure mean match are

$$\begin{aligned} \mathbb{E}_d[R_i \mid \widehat{\zeta}_i, u] &= 1, \\ \mathbb{E}_d[R_i Y_i \mid \widehat{\zeta}_i, u] &= \widehat{\zeta}_i e^{-b_i u} + \vartheta_i^S u \phi(b_i u), \quad \vartheta_i^S = \vartheta_i + \mathfrak{I}_i \sqrt{\boldsymbol{\Psi}_i}. \end{aligned} \quad (4.232)$$

The deterministic maps leave the stock unchanged and preserve the first identity. The second identity does not assert an exact match of $\mathbb{E}_d[R_i \zeta_i^{\text{next}}]$ or $\mathbb{E}_d[R_i \xi_i^{\text{next}}]$ to the full nonlinear stock-measure transition. The retained exponential endpoint tilt matches a CIR first moment, not its entire changed-measure law. These nonlinear

moments and finite-step numeraire consistency must be checked numerically. The correlation $\mathfrak{N}_i$ retains its signed ζ_i -stock meaning; in ξ_i coordinates the corresponding correlation gains $\text{sign}(\lambda_i)$.

The residual rough log variance is approximated by

$$(1 - \mathfrak{N}_i^2) \frac{\mathfrak{V}_i}{2} \mathbb{E}_d \left[\int_0^u \frac{dv}{Y_i(v)} \mid Y_i(0) = \widehat{\zeta}_i \right], \quad (4.233)$$

where the expectation is for the intermediate affine CIR block. The existing reciprocal fair-variance formula and its 24-node quadrature evaluate this integral. In the corrected branch this is an approximation to the integrated variance of the nonlinear rough process; it is not an exact nonlinear fair-variance or conditional bridge-variance formula. The smooth and rough conditional stock factors multiply and their residual variances accumulate. All stock normalizations remain analytic, with no simulated sample mean.

This is a finite-step composition, not an exact transition of the coupled model. At $\gamma_i = 0$, the rough block is affine, but its composition in ζ_i with the smooth transition in ξ_i generally differs from the former merged-Gamma update in ξ_i . The corrected branch retains composition error; other nonlinear branches retain frozen-rate error. Earlier confirmations from both older schemes are historical results. Their parameters may initialize the corrected scheme, but their errors must be recomputed under the new numerical revision. Outer-step convergence, transformed stock-measure moments and integrated-variance accuracy remain to be measured at the retained parameters.

Active singular or nonrepresentable values are reported without clipping or epsilon floors. Native Torch sampling, full-path vectorization and the AMD/ROCm backend are retained without path batching. Direct residual stock sampling remains once per selected expiration.

Calibration continuation. The new implementation, notebook, checkpoints and results preserve the earlier variant. There are 24 model coefficients, of which 22 are free if the two already near-unit correlations retain their existing fixed status. The 37 expiry speeds remain a date dictionary with the first fixed at one. The target remains 2026-10-05 EOD reported close and the same 37 expirations through 2027-02-26, within the explicit maturity limit 2027-03-12. Both rights and all positive-weight ITM and OTM strikes remain in the volume/bid/ask price-space objective. The existing staged direct-pricing grids and path counts, conditional-pricing confirmations, raw forward errors and their standard errors are retained. Ongoing confirmations use the configured 2^{19} paths; later final convergence uses at least 2^{20} . No calibration improvement or grid convergence is asserted before the corresponding measurements.

Three power-coordinate states with cyclic rough clocks

For $i = 1, 2, 3$, define $j(i) = 1 + (i \bmod 3)$ and

$$\zeta_i = \mathfrak{I}_i \xi_i^{\lambda_i}, \quad \mathfrak{I}_i > 0, \quad \lambda_i \neq 0. \quad (4.234)$$

There are three state variables in two coordinate representations. The rough block of ζ_1 uses $\mathfrak{r}_2$, that of ζ_2 uses $\mathfrak{r}_3$, and that of ζ_3 uses $\mathfrak{r}_1$. The orientation of this directed cycle is fixed. There is no magic clock.

Process, clocks and covariances. The assigned-clock affine components under the domestic measure are

$$\begin{aligned} \mathbb{E}_d[d\xi_i]_{dt} &= (\theta_i - \kappa_i \xi_i)dt, & d\langle \xi_i \rangle_{dt} &= 2\xi_i dt, \\ \mathbb{E}_d[d\zeta_i]_{d\mathfrak{r}_{j(i)}} &= (\vartheta_i - \mathfrak{I}_i \zeta_i)d\mathfrak{r}_{j(i)}, & d\langle \zeta_i \rangle_{d\mathfrak{r}_{j(i)}} &= 2\zeta_i d\mathfrak{r}_{j(i)}. \end{aligned} \quad (4.235)$$

The state martingales are orthogonal: $d\langle \xi_i, \xi_k \rangle = 0$ for $i \neq k$. The discounted stock has $\mathbb{E}_d[dS] = 0$ and

$$\begin{aligned} \frac{d\langle S \rangle}{S^2} &= \frac{1}{2} \sum_{i=1}^3 \left[H_i \xi_i dt + \frac{\mathfrak{P}_i}{\zeta_i} d\mathfrak{r}_{j(i)} \right], \\ \frac{d\langle \xi_i, S \rangle}{S} \Big|_{dt} &= \rho_i \sqrt{H_i} \xi_i dt, \\ \frac{d\langle \zeta_i, S \rangle}{S} \Big|_{d\mathfrak{r}_{j(i)}} &= \mathfrak{N}_i \sqrt{\mathfrak{P}_i} d\mathfrak{r}_{j(i)}. \end{aligned} \quad (4.236)$$

All coefficients are state indexed. The component Brownian correlations obey $|\rho_i|, |\mathfrak{N}_i| \leq 1$; in the rough ξ_i coordinate the correlation is $\text{sign}(\lambda_i)\mathfrak{N}_i$. Independent residual stock channels supply the remaining variance.

Each state generates its own boundary clock with the same normalization:

$$\mathfrak{r}_i(t) = \lim_{\epsilon \downarrow 0} \frac{\theta_i(1 - \theta_i)}{\epsilon^{\theta_i}} \int_0^t \mathbf{1}_{\{\xi_i(s) < \epsilon\}} ds, \quad \mathbf{1}_{\{\xi_i > 0\}} d\mathfrak{r}_i = 0. \quad (4.237)$$

The inherited domain requires $0 < \theta_i < 1$ and $\theta_i + \theta_k > 1$ for every distinct pair, with nonnegative initial states and no initially zero pair. A condition on $\theta_1 + \theta_2 + \theta_3$ alone does not control pairwise corners. The pairwise restrictions are the intended boundary safeguard; their sufficiency for the fully coupled process has not been proved. Distinct contacts within a finite numerical interval are not simultaneous zeros. For positive λ_i , the rough stock coefficient diverges as $\xi_i \downarrow 0$; its intended protection is inactivity of the next state's clock, subject to this unresolved boundary qualification.

On the positive interior the complete smooth-coordinate equations are

$$\begin{aligned}
\mathbb{E}_d[d\xi_i] &= (\theta_i - \kappa_i \xi_i)dt + \left[\frac{\vartheta_i - 1 + 1/\lambda_i}{\lambda_i \mathfrak{D}_i} \xi_i^{1-\lambda_i} - \frac{\mathfrak{J}_i}{\lambda_i} \xi_i \right] d\mathfrak{r}_{j(i)}, \\
d\langle \xi_i \rangle &= 2\xi_i dt + \frac{2\xi_i^{2-\lambda_i}}{\lambda_i^2 \mathfrak{D}_i} d\mathfrak{r}_{j(i)}, \\
\frac{d\langle \xi_i, S \rangle}{S} &= \rho_i \sqrt{H_i} \xi_i dt + \frac{\mathfrak{N}_i \sqrt{\Psi_i}}{\lambda_i \mathfrak{D}_i} \xi_i^{1-\lambda_i} d\mathfrak{r}_{j(i)}.
\end{aligned} \tag{4.238}$$

Adding the state–stock covariance preserves the local family under the stock numeraire:

$$\kappa_i^S = \kappa_i - \rho_i \sqrt{H_i}, \quad \vartheta_i^S = \vartheta_i + \mathfrak{N}_i \sqrt{\Psi_i}. \tag{4.239}$$

The other coefficients and clock assignment are unchanged. Inverting the stock quote reverses both channel correlations. The retained conditions $\vartheta_i > 1$ and $\vartheta_i^S > 1$ give a finite inverse first moment for each isolated affine rough block. These local conditions do not establish global existence or true martingality of the coupled model.

Parameter names and dimensions. For each $i = 1, 2, 3$, the eleven descriptor names correspond to

symbol	descriptor name	time dimension
θ_i	theta_i	1
κ_i	kappa_i	T^{-1}
$\xi_{i,0}$	xi_i_0	T
H_i	H_i	T^{-2}
$\mathfrak{D}_i$	kaf_i	$T^{1-\theta_{j(i)}-\lambda_i}$
λ_i	lambda_i	1
ϑ_i	vartheta_i	1
$\mathfrak{J}_i$	nun_i	$T^{\theta_{j(i)}-1}$
Ψ_i	evar_i	1
ρ_i	rho_i	1
$\mathfrak{N}_i$	ecorr_i	1

Thus $[\mathfrak{r}_{j(i)}] = [\zeta_i] = T^{1-\theta_{j(i)}}$. The calibration independently fits all 33 model coefficients. Explicit `param(initial, change scale, bounds)` entries retain the mathematical bounds, including sign-preserving half-lines for nonzero λ_i , without a finite magnitude cap. The five expiry speeds are a nested date dictionary initialized at one; the first is fixed and the remaining four are free.

Cyclic sign classes and symmetric initial values. Cyclic rotations produce four sign classes, represented by $(- - -)$, $(+ + +)$, $(- - +)$ and $(+ + -)$, named **nnn**, **ppp**, **npn** and **ppn**. All state-indexed coefficients and initial values must rotate together. Reflections reverse the directed cycle and are unnecessary for this classification. Each representative has its own descriptors, checkpoints, statistics and plots. The symmetric starting values are initial guesses; no equality constraints are imposed during fitting.

The initial negative-power seed is constructed from the saved first-week two-state **nn** parameters, not from a purported exact embedding. Writing the old initial states and coordinates as x_a, z_a , $a = 1, 2$, set

$$\begin{aligned}\bar{x} &= \sqrt{x_1 x_2}, & \bar{z} &= \sqrt{z_1 z_2}, & \bar{\mathfrak{z}} &= \bar{z} \bar{x}^{-\bar{\lambda}}, \\ \bar{H} &= \frac{H_1 x_1 + H_2 x_2}{3\bar{x}}, & \bar{\Psi} &= \frac{\bar{z}}{3} \left(\frac{\Psi_1}{z_1} + \frac{\Psi_2}{z_2} \right).\end{aligned}\tag{4.240}$$

The remaining seed coefficients $\theta, \kappa, \lambda, \vartheta, \mathfrak{z}, \rho, \mathfrak{D}$ are arithmetic means of the old values. These shared values are copied into all three states. This construction preserves the summed initial smooth variance and the summed initial rough stock coefficients, not complete paths or prices. Positive signs use the existing local map (4.226). The exact source snapshot, its hash and constructed values are retained in `cyclic_power_reciprocal_variance_initializer.json`. Zeroing state 3's stock weights would still leave its clock driving state 2, so it would not recover the previous two-state cycle.

Unchanged single-state numerical procedure. The separate three-state kernel calls the existing single-state coefficient, merged-transition and stock-step functions without altering their formulas. For each outer interval it first draws all three physical contact records from the same pre-step state vector, at the individual rates κ_i . State i then receives record $j(i)$'s rough increment. There are exactly three merged endpoints, one for each state, before the old states are replaced.

In particular, at interval-start state $x = \xi_i$, the frozen quantities are precisely those of (4.227). For $\lambda_i > 0$ use $a_i = q_i[1 + \lambda_i(\vartheta_i - 1)]$ and $b_i = \mathfrak{z}_i/\lambda_i$; for $\lambda_i < 0$ use $a_i = 0$ and $b_i = \mathfrak{z}_i/\lambda_i - q_i[1 + \lambda_i(\vartheta_i - 1)]/x$. For physical conditional inputs (c_i, s_i, r_i) the endpoint retains

$$\begin{aligned}D_i &= s_i + q_i \delta \mathfrak{r}_{j(i)}, & A_i &= c_i s_i + a_i \delta \mathfrak{r}_{j(i)}, & B_i &= b_i \delta \mathfrak{r}_{j(i)}, \\ \text{shape}_i &= A_i/D_i, & \text{scale}_i &= D_i \phi(B_i), & \text{noncentral}_i &= r_i e^{-B_i},\end{aligned}\tag{4.241}$$

where $\phi(B) = (1 - e^{-B})/B$ with $\phi(0) = 1$. Positive powers retain the rough immigration in the Gamma shape. The original stock-measure mean match uses ϑ_i^S in that immigration; for negative powers it uses the original proportional-rate

shift. The analytic physical-clock likelihood, normalized Gamma endpoint tilt and covariance-residual projection are unchanged. The three log-forward and residual-variance contributions are added.

This remains the established finite-step approximation. There is one merged update per configured outer interval, with no operator split or extra subdivision. Inactive rough increments do not evaluate singular powers; active nonrepresentable coefficients fail explicitly without clipping. Residual stock normals are sampled once per selected expiry and carried forward. Direct option payoffs retain Torch sorting and cumulative sums; conditional integration remains the confirmation estimator. No empirical centering, price rescaling or sample correction is introduced.

First five expirations and validation scope. The market target is the 2026-10-05 reported close, with the first five available expirations 2026-10-06, 07, 08, 09 and 12 and explicit maturity limit 2026-10-12. Both calls and puts, ITM and OTM, at all supplied positive-volume strikes enter the existing volume/bid/ask objective. Fitting retains direct pricing at 2^{15} , 2^{17} and 2^{19} paths and the existing early quarter-day grid, with maximum intervals 5, 2 and 1 days. Conditional confirmations use 2^{19} paths and at-most-daily intervals. The later final time-step convergence study retains its separate 2^{20} minimum.

The four signs run sequentially on the AMD GPU through the shared Jupyter server. Every branch resumes its own saved best and retains the full fitting-statistics display, expiry plots and raw forward errors with their Monte Carlo standard errors. The earlier two-state experiment is paused with its results preserved separately. CPU checks cover the cyclic record assignment, the generator and stock-numeraire identities, exact reuse of the single-state formulas, endpoint counts and raw forward behavior for all four sign representatives. This specification does not assert a completed fit or a time-step convergence result.

Magic time and an affine physical clock

The construction in (4.358) can be applied to the complete independent reciprocal specification, including the optional Heston contribution in (4.112). Let τ denote process time, or *magic time*. All existing internal dynamics run in τ : replace dt by $d\tau$, generate the rough clock $\mathfrak{r}_\tau^\xi$ from ξ_τ , and retain the existing stock and reciprocal-state coefficients. A separate affine process determines when this system is observed in physical time:

$$d\chi_\tau = (\theta_\chi - \kappa_\chi \chi_\tau) d\tau + \sqrt{2\chi_\tau} dW_\tau^\chi, \quad \chi_0 = 0, \quad 0 < \theta_\chi < 1, \quad \kappa_\chi > 0. \quad (4.242)$$

The subscript χ identifies this clock driver. It replaces a long descriptive time subscript without changing the meanings of the existing θ and κ parameters. In the

first implementation, χ is independent of all existing state and stock innovations.

Let $\mathfrak{r}_\tau^\chi$ be its canonical boundary clock, with the same normalization as the earlier affine boundary clocks. Define the physical clock and its right-continuous inverse by

$$C_\tau = \frac{\Gamma(\theta_\chi)}{(1 - \theta_\chi)\kappa_\chi^{\theta_\chi}} \mathfrak{r}_\tau^\chi, \quad \tau(u) = \inf\{s : C_s > u\}. \quad (4.243)$$

Here C , rather than χ , is increasing. It is constant during an excursion of χ away from zero, and $\tau(u)$ skips the entire excursion. Observed stock and volatility states therefore jump together, although their trajectories in magic time are continuous. The observing driver is at the boundary at either end of each skipped excursion. In particular, every simulation started from an observed valuation date starts this driver at $\chi = 0$, while retaining the observed values of ξ and ζ . The boundary-start inverse-clock law below applies to each such future simulation; there is no additional initial observing state to calibrate.

For a calendar expiration T , let $u(T)$ denote the existing deterministic holiday-adjusted and fitted time map. The observed states are

$$S_T = S_{\tau(u(T))}, \quad \xi_T = \xi_{\tau(u(T))}, \quad \zeta_T = \zeta_{\tau(u(T))}. \quad (4.244)$$

Thus $u(T)$ is the deterministic holiday-adjusted and fitted physical-time map, applied once, while $\tau(u(T))$ is the random magic-time horizon. Market expirations, forwards, discounting, and implied-volatility annualization retain their calendar conventions. No additional fitted speed per expiration is introduced by the stochastic clock itself. When the deterministic time map is recalibrated, its positive interval speeds act on holiday-adjusted physical time before the inverse clock is sampled. The order is calendar time, holiday-adjusted time, $u(T)$, and then $\tau(u(T))$. Multiplying the sampled magic-time horizons by an expiration-dependent factor would specify a different construction.

Inverse-clock law and the previous-model limit. For $q \geq 0$, the independent increments of the normalized inverse clock satisfy

$$\begin{aligned} \mathbb{E}[e^{-q\Delta\tau}] &= \exp\{-\Delta u \Phi_\chi(q)\}, \\ \Phi_\chi(q) &= \kappa_\chi \frac{\Gamma(q/\kappa_\chi + 1 - \theta_\chi)}{\Gamma(q/\kappa_\chi)\Gamma(1 - \theta_\chi)}. \end{aligned} \quad (4.245)$$

The value at $q = 0$ is defined by continuity. Differentiation at zero gives

$$\mathbb{E}[\Delta\tau] = \Delta u, \quad \text{Var}(\Delta\tau) = \frac{2\Delta u}{\kappa_\chi} [\psi(1) - \psi(1 - \theta_\chi)], \quad (4.246)$$

where ψ is the digamma function. Mean reversion makes both moments finite; the mean-reversion-free inverse boundary clock instead has a stable law with infinite first moment. The normalization in (4.243) fixes the mean magic-time rate at one. As $\theta_\chi \downarrow 0$, $\Phi_\chi(q) \rightarrow q$ and $\tau(u) \rightarrow u$ in mean square at every fixed u . The implementation defines $\theta_\chi = 0$ as an exact deterministic-clock branch, recovering the previous simulator and its prices without evaluating the singular boundary-clock normalization. Starting χ away from zero would add an initial hitting delay and would not have this nesting property.

The extended calibration therefore includes the previously fitted model, with all its financial coefficients and its deterministic physical-time map unchanged. For the same market target, the infimum of the extended calibration error cannot exceed the previous-model infimum. A finite numerical search need not find that infimum, so the previous fit is retained as an explicit candidate at $\theta_\chi = 0$. Comparisons use the same path count, simulation grid, quadrature, and random seeds; scores obtained at different resolutions do not by themselves establish an improvement. A positive-clock candidate is promoted only after a comparison with that retained fit at common settings and a subsequent refinement check. The observing-clock coefficients can first be fitted on the shortest expirations, then held fixed while extending the calibration to longer expirations. Keeping θ_χ relatively small helps localize the additional clock dispersion: by (4.246), $\text{Var}(K_U)/U^2$ decreases as $1/U$. This does not ensure that the longer option smiles remain unchanged, so those expirations are repriced and compared after the short-end fit.

Both θ_χ and the original affine Feller characteristics are dimensionless. The quantities χ , τ , C , and u have dimension time, while κ_χ has dimension time^{-1} . The canonical clock $\mathfrak{r}^\chi$ has dimension $\text{time}^{1-\theta_\chi}$, so its multiplier in (4.243) has dimension $\text{time}^{\theta_\chi}$. Changing θ_χ therefore does not change the units of either physical or magic time.

The scaling rules for the base coefficients nevertheless refer to magic time, not to the typical short physical-time horizon. For $0 < \theta_\chi < 1$, the large- q behavior of (4.245) is

$$\Phi_\chi(q) \sim \frac{\kappa_\chi^{\theta_\chi}}{\Gamma(1-\theta_\chi)} q^{1-\theta_\chi}.$$

Consequently a fixed interior quantile of K_U , as $U \downarrow 0$, scales as a constant times

$$\left(\frac{\kappa_\chi^{\theta_\chi} U}{\Gamma(1-\theta_\chi)} \right)^{1/(1-\theta_\chi)}.$$

This is a distributional scaling, not a scaling of the mean: $\mathbb{E}[K_U] = U$ still holds. The homogeneous base-model powers $\xi, \zeta \sim \tau$ and $\mathfrak{r}^\xi \sim \tau^{1-\theta}$ therefore compose with this physical-time power. They do not imply that increments from a strictly

positive initial state obey the same boundary scaling. At long physical horizons, $K_U/U \rightarrow 1$ in mean square, so the linear time interpretation is recovered.

An existing characteristic magic-time duration $\Delta\tau$ can instead be assigned an exact physical reference duration by

$$U(\Delta\tau) = \frac{1}{\Phi_\chi(1/\Delta\tau)}, \quad \mathbb{E}[e^{-K_U(\Delta\tau)/\Delta\tau}] = e^{-1}. \quad (4.247)$$

For example, the physical relaxation time of an affine mean with magic mean-reversion rate κ is $1/\Phi_\chi(\kappa)$, rather than $1/\kappa$. At $\theta_\chi = 0$ these durations coincide. This convention keeps intrinsic magic-time units separate from the effect of the observation clock and from the deterministic fitted map $u(T)$.

Simulation without resolving the observing driver. Independence allows the random magic-time horizons to be generated before the existing state paths. The Lévy measure of the inverse clock has tail

$$\bar{\nu}_\chi(s) = \nu_\chi((s, \infty)) = \frac{\kappa_\chi}{\Gamma(1 - \theta_\chi)\Gamma(\theta_\chi)} (e^{\kappa_\chi s} - 1)^{-(1-\theta_\chi)}, \quad s > 0. \quad (4.248)$$

Choose a positive resolution ϵ in magic-time units. On an interval of physical-clock length Δu , sample a Poisson number of resolved jumps with mean $\Delta u \bar{\nu}_\chi(\epsilon)$. Their sizes have survival probability $\bar{\nu}_\chi(s)/\bar{\nu}_\chi(\epsilon)$ for $s \geq \epsilon$, so inversion uses only exponentials and logarithms. Replace unresolved jumps by their deterministic mean, giving

$$\Delta\tau_\epsilon = \Delta u b_\epsilon + \sum_{j=1}^N J_j, \quad b_\epsilon = \int_0^\epsilon s \nu_\chi(ds), \quad N \sim \text{Poisson}(\Delta u \bar{\nu}_\chi(\epsilon)). \quad (4.249)$$

This approximation preserves $\mathbb{E}[\Delta\tau_\epsilon] = \Delta u$ exactly. Its missing clock variance is

$$\begin{aligned} v_\epsilon &= \int_0^\epsilon s^2 \nu_\chi(ds) = 2 \int_0^\epsilon s \bar{\nu}_\chi(s) ds - \epsilon^2 \bar{\nu}_\chi(\epsilon), \\ 0 \leq v_\epsilon &\leq \frac{2\kappa_\chi^{\theta_\chi} \epsilon^{1+\theta_\chi}}{(1 + \theta_\chi)\Gamma(1 - \theta_\chi)\Gamma(\theta_\chi)}. \end{aligned} \quad (4.250)$$

Under the natural coupling with the exact jump clock, $\mathbb{E}[(\Delta\tau_\epsilon - \Delta\tau)^2] = \Delta u v_\epsilon$. The finite-resolution clock is not an exact sample of (4.245); reducing ϵ removes this approximation. The separate state-grid and effective-duration refinements of the existing simulator remain necessary. At a fixed positive resolution, deterministic small-jump compensation does not reproduce the stable typical-horizon scaling all the way to $\Delta u = 0$. That scaling requires the cutoff to become small relative to

the typical magic-time increment. Short-expiration calibration must therefore be checked by reducing the clock cutoff as well as the state step size.

For each path, accumulate these nonnegative magic-time increments at the requested physical expirations and advance the complete existing system to its successive random horizons. All state changes and both internal stock-variance channels evolve over a skipped magic-time interval. The physical-time process must not be constructed by linearly interpolating its states through that interval: the observing clock assigns the full change to a jump. The new clock is independent of the base process, so a true stock martingale in magic time remains a martingale after this time change. Conditional discrete stock normalizers can likewise be retained on every random-duration state step. This argument does not cover a clock driver correlated with stock innovations; such a specification would require a separate analysis of the stopped stock and its integrability.

VIX after the jump time change. For a calendar window $\Delta_c = 30/365$, the independent future magic-time duration has Laplace transform (4.245) with $\Delta u = u(T + \Delta_c) - u(T)$. Conditional on the observed base states, average the magic-time expected log-contract payoff over this duration distribution before taking its square root:

$$\text{VIX}_T^{\log} = 100 \sqrt{\frac{-2}{\Delta_c} \mathbb{E}_d \left[\log \frac{S_{\tau(u(T + \Delta_c))}}{S_{\tau(u(T))}} \middle| \xi_T, \zeta_T \right]}. \quad (4.251)$$

The stock is in the discounted convention used above. For the continuous base model, its expected log-contract payoff at a given magic-time duration equals its expected integrated total stock variance, including any Heston contribution. Consequently the reciprocal fair-variance formula remains useful conditional on duration, and the existing rough-clock approximation remains an additional approximation inside that conditional calculation. Substituting $\mathbb{E}[\Delta\tau]$ for the duration is not generally exact.

After the time change, sums of squared observed log-price jumps are not path-wise equal to the integrated variance in magic time. The log-contract target in (4.251) must therefore be distinguished from observed jump quadratic variation. This section specifies the extension and its sampler; it does not report a completed calibration or establish an exact closed-form VIX approximation for the full rough model.

Conditional fair variance in magic time

Let K_U denote the future magic-time increment corresponding to a physical-clock interval of length U , with transform (4.245). For the continuous base system, define

the conditional fair area and its independent-clock average by

$$\begin{aligned}
G_{\text{base}}(\tau; \xi, \zeta) &= \mathbb{E}_{\text{d}} \left[\int_0^\tau \left(\frac{\mathcal{H}}{2} \xi_v + \frac{1}{\zeta_v} \right) dv \right. \\
&\quad \left. + \int_0^\tau \frac{\Psi_S}{\zeta_v} d\mathfrak{K}_v^\xi \mid \xi_0 = \xi, \zeta_0 = \zeta \right], \\
G_\chi(U; \xi, \zeta) &= \mathbb{E} [G_{\text{base}}(K_U; \xi, \zeta)].
\end{aligned} \tag{4.252}$$

The second equality is exact whenever the expected area is finite. It separates the existing state dynamics from the new, independent observation clock. For a calendar window, use $U = u(T + \Delta_c) - u(T)$ in (4.252). Equation (4.251) then gives

$$\text{VIX}_T^{\log} = 100 \sqrt{G_\chi(U; \xi_T, \zeta_T) / \Delta_c}.$$

Average fair areas before taking the square root and retain the calendar annualization denominator. This remains the expected log-contract quantity, not the squared-jump quadratic variation of the observed stock.

Exact Heston contribution. Linearity of expected area permits the Heston and reciprocal channels to be evaluated separately even though their state dynamics can be linked. For $\kappa > 0$, averaging the affine mean of ξ gives

$$G_{\chi, \xi}(U; \xi) = \frac{\mathcal{H}}{2} \left[\frac{\theta}{\kappa} U + \left(\xi - \frac{\theta}{\kappa} \right) \frac{1 - e^{-U \Phi_\chi(\kappa)}}{\kappa} \right]. \tag{4.253}$$

At $\kappa = 0$, its continuous limit uses the second moment of the magic horizon:

$$G_{\chi, \xi}(U; \xi) = \frac{\mathcal{H}}{2} \left[\xi U + \frac{\theta}{2} (U^2 + \text{Var}(K_U)) \right]. \tag{4.254}$$

These formulas include the total Heston variance; no orthogonal-variance factor such as $1 - \rho_\xi^2$ appears.

Exact smooth reciprocal contribution. For the pure smooth reciprocal channel, set its rough state and stock coefficients to zero and write $\nu = 1 + 2\alpha/\Psi_{\zeta, t}$ and $c_\infty = \Psi_{\zeta, t}/(2\beta)$, with $\Psi_{\zeta, t} > 0$, $\alpha > 0$, $\beta > 0$, and $\zeta > 0$. The Itô log-state identity gives its fair area as $[\mathbb{E} \log \zeta_\tau - \log \zeta + \beta\tau]/\alpha$. Expanding the log state in

the CIR Laguerre eigenfunctions and averaging their exponential time factors gives

$$G_{\chi,\zeta}(U; \zeta) = \frac{1}{\alpha} \left[\beta U + \log \frac{c_\infty}{\zeta} + \psi(\nu) - \Gamma(\nu) \sum_{n=1}^{\infty} \frac{n! L_n^{(\nu-1)}(\zeta/c_\infty)}{n \Gamma(\nu+n)} e^{-U \Phi_\chi(n\beta)} \right]. \quad (4.255)$$

Here $L_n^{(\nu-1)}$ is the generalized Laguerre polynomial. The log-state coefficient follows from $\int_0^\infty x^{\nu-1} e^{-x} L_n^{(\nu-1)}(x) \log x \, dx = -\Gamma(\nu)/n$ for $n \geq 1$. For $U > 0$ the series is convergent, but its numerical truncation must be checked. Since $\Phi_\chi(n\beta)$ grows proportionally to $n^{1-\theta_\chi}$, damping can be weak at short horizons or when θ_χ is close to one. Large ζ/c_∞ can also cause severe cancellation among polynomial terms. Compare increasing mode counts and reject a spectral evaluation that does not meet its numerical checks; select positive Gamma-horizon quadrature or a sampled-horizon estimate explicitly rather than silently changing the horizon law. The implementation anchors the correction to the positive fixed-duration fair-area evaluator of (4.123):

$$G_{\chi,\zeta}(U; \zeta) = G_{\text{base},\zeta}(U; \zeta) - \frac{\Gamma(\nu)}{\alpha} \sum_{n=1}^{\infty} \frac{n! L_n^{(\nu-1)}(\zeta/c_\infty)}{n \Gamma(\nu+n)} \left[e^{-U \Phi_\chi(n\beta)} - e^{-Un\beta} \right]. \quad (4.256)$$

At $\theta_\chi = 0$ the correction vanishes exactly. The fixed-state Laguerre asymptotics are given in NIST's Digital Library of Mathematical Functions.⁶

Positive approximate evaluation for the rough channel. When the existing rough-state fair area is available only through a conditional mean-clock or frozen-clock approximation, denote that nonnegative base estimate by $\hat{G}_{\text{base},\zeta}$. An inexpensive outer approximation matches the first two moments of the independent magic horizon:

$$K_U \approx \text{Gamma} \left(\text{shape} = \frac{U^2}{\text{Var}(K_U)}, \quad \text{scale} = \frac{\text{Var}(K_U)}{U} \right), \quad U > 0. \quad (4.257)$$

Positive quadrature then averages $\hat{G}_{\text{base},\zeta}(K_U; \xi, \zeta)$ under this Gamma law, while retaining (4.253) for the Heston part. At $\theta_\chi = 0$, use $K_U = U$ directly, recovering the chosen base evaluator without Gamma parameters or horizon quadrature; zero horizon always has zero area.

⁶Equation (18.15.14): <https://dlmf.nist.gov/18.15.E14>.

There are two distinct approximations here: replacing the inner rough clock by its conditional mean, and replacing the outer magic horizon by a Gamma law. Neither is present in the exact identity (4.252), and the pure smooth spectral benchmark eliminates the outer approximation. Moment matching alone does not establish accuracy in the tails. In particular, at fixed κ_χ the exact clock's third cumulant is of first order in θ_χ near zero, whereas the matched Gamma third cumulant is of second order. Small θ_χ alone therefore does not justify this closure or a variance-only Taylor correction. The curvature of the base fair area can change sign with state and horizon, so the mean-horizon substitution has no universal upper- or lower-bound interpretation. Check horizon approximations against the exact smooth benchmark and, for the full rough system, against independently sampled horizons and future state paths; quadrature refinement controls only the numerical integration error of the selected approximation. The sampled-horizon alternative can use the native compensated-jump clock sampler with its reported cutoff, a linear control variate based on $\mathbb{E}[K_U] = U$, and a Monte Carlo standard error. It retains the clock-resolution and inner base-approximation errors, separately from its sampling error.

Reciprocal volatility with magic and physical clocks

A separate experiment removes the volatility boundary clock altogether. Both the stock and its reciprocal variance state retain contributions from magic time τ and physical time t . Only the separate rough volatility clock is removed. An excursion of the physical-clock driver is skipped when the system is observed, so the stock and volatility can jump together. This tests whether a strong short-expiry smile can be explained by the observation clock without the large dispersion of conditional future variance found in some of the earlier rough-clock calibrations. It is a modeling hypothesis, not an empirical identification of roughness with jumps. The observed process has jumps, and its finite-frequency irregularity must not be identified with the local Hurst exponent of a continuous rough-volatility model.

Specification and observable state. In the discounted-stock convention, the complete magic-time moments are

$$\begin{aligned}
\mathbb{E}_d[dS_\tau] &= 0, & d\langle S \rangle_\tau &= \frac{S_\tau^2}{\zeta_\tau}(d\tau + \mathfrak{v}_{S,t} dt), \\
\mathbb{E}_d[d\zeta_\tau] &= \left(\frac{\vartheta}{2}\mathfrak{v}_\zeta - \beta\zeta_\tau\right) d\tau + \left(\frac{\mathfrak{v}_{\zeta,t}}{2} + \alpha_t - \beta_t\zeta_\tau\right) dt, & d\langle \zeta \rangle_\tau &= \zeta_\tau(\mathfrak{v}_\zeta d\tau + \mathfrak{v}_{\zeta,t} dt), \\
d\langle \zeta, S \rangle_\tau &= -S_\tau(\rho\sqrt{\mathfrak{v}_\zeta} d\tau + \rho_t\sqrt{\mathfrak{v}_{\zeta,t}\mathfrak{v}_{S,t}} dt), \\
\mathbb{E}[d\chi_\tau] &= (\theta_\chi - \kappa_\chi\chi_\tau)d\tau, & d\langle \chi \rangle_\tau &= 2\chi_\tau d\tau, \\
d\langle \chi, \zeta \rangle_\tau &= d\langle \chi, S \rangle_\tau = 0, & \chi_0 &= 0.
\end{aligned} \tag{4.258}$$

Here ρ is the instantaneous variance–stock correlation; the reciprocal-state–stock correlation is $-\rho$ in magic time and $-\rho_t$ in physical time. A negative ρ therefore gives the intended negative volatility–stock leverage. Fixing the stock magic-time variance coefficient at one retains the existing scale of ζ . The earlier rough-volatility driver ξ and its boundary clock are absent; there are no coefficients multiplying its rough clock. The five physical coefficients are retained: stock strength $\mathfrak{v}_{S,t}$, state strength $\mathfrak{v}_{\zeta,t}$, excess constant drift α_t , proportional drift β_t , and volatility–stock correlation ρ_t . With active physical state and stock diffusion, finite fair variance requires $\alpha_t > 0$, giving physical cash Feller characteristic $1 + 2\alpha_t/\mathfrak{v}_{\zeta,t} > 1$. In stock units it becomes $1 + 2(\alpha_t - \rho_t\sqrt{\mathfrak{v}_{\zeta,t}\mathfrak{v}_{S,t}})/\mathfrak{v}_{\zeta,t}$. The physical clock does not advance during a skipped magic excursion.

Use the normalized physical clock and inverse already defined in (4.243), with the existing Hebrew symbol $\mathfrak{x}$ for its boundary clock. If $u(T)$ is the deterministic holiday-adjusted physical-time map, observations are

$$S_T = S_{\tau(u(T))}, \quad \zeta_T = \zeta_{\tau(u(T))}, \quad \chi_{\tau(u(T))} = 0. \tag{4.259}$$

The physical-clock driver is always initialized at zero at an observed valuation date. Only ζ_T remains as an observable volatility state; there is no independent initial physical-driver state to fit. The independent future magic duration K_U retains the transform (4.245), with $U = u(T + \Delta_c) - u(T)$. These transform identities describe the exact boundary clock. They provide a benchmark for the regular-physical-grid surrogate constructed below. Omitting the launch endpoint noise changes its clock law; grid refinement alone is not claimed to recover the exact inverse-clock law. The deterministic calendar-to-physical map is applied before sampling the magic horizon. It is not applied to an already sampled magic duration. At $\theta_\chi = 0$, the physical clock is deterministic and the model reduces exactly to the ordinary smooth reciprocal specification.

The model parameters are $\zeta_0, \vartheta, \beta, \mathfrak{v}_\zeta, \rho, \theta_\chi, \kappa_\chi, \mathfrak{v}_{S,t}, \mathfrak{v}_{\zeta,t}, \alpha_t, \beta_t, \rho_t$, together with any explicitly configured deterministic physical-time map. With the normalized

mean clock, ζ_0 has the dimension of magic time; β and κ_χ have inverse-magic-time dimension; and the remaining listed parameters are dimensionless. Their interpretation at short physical horizons also depends on θ_χ , as described after (4.247).

Admissibility. For stochastic reciprocal volatility, require $\zeta_0 > 0$, $\Psi_\zeta > 0$, $|\rho| \leq 1$, and $\vartheta > 1$. The last condition is required for finite conditional expected inverse state and fair variance. At $\vartheta = 1$, nonattainment of zero is not enough: the inverse-state expectation, and hence the fair variance, diverges. Approaching this value can also make finite-grid variance estimation unreliable; this is a reason for explicit numerical checks, not an arbitrary universal margin such as a fixed characteristic of 1.5. Under the stock numeraire the characteristic is

$$\vartheta_f = \vartheta - \frac{2\rho}{\sqrt{\Psi_\zeta}}. \quad (4.260)$$

Its boundary nonattainment condition must hold as well. For the intended negative leverage, $\vartheta_f \geq \vartheta$ automatically. The observation clock requires $0 \leq \theta_\chi < 1$ and $\kappa_\chi > 0$; the zero-characteristic case is handled as the exact deterministic branch. There is no mathematical sign restriction on β for finite-horizon state transitions. Any narrower numerical domain, finite-moment requirement of a stock-factor approximation, or optimizer limitation must be stated in the calibration configuration. The retained endpoint-bridge implementation currently supports $-1 \leq \rho \leq 0$; this implements the intended negative leverage and does not establish that positive leverage is inadmissible in the continuous model.

Simulation and calibration. The inverse-clock transform is available for nonzero κ_χ ; this does not imply an elementary density or a closed-form quantile. One way to check its normalization without using an inverse-clock sampler is through the boundary potential. The Gamma transition law from $\chi_0 = 0$ gives the normalized boundary-clock mean rate $(1 - e^{-\kappa_\chi s})^{-\theta_\chi}$. Consequently,

$$\int_0^\infty e^{-qs} d\mathbb{E}[C_s] = \frac{1}{\kappa_\chi} B(q/\kappa_\chi, 1 - \theta_\chi) = \frac{1}{\Phi_\chi(q)}, \quad q > 0.$$

Here B is the beta function. The last equality follows by changing variables from the continuous boundary clock to its inverse: $\int e^{-qs} dC_s = \int e^{-q\tau(u)} du$. This recovers (4.245), including its mean-reversion dependence, and agrees with the canonical transform in (4.53).

At a deterministic physical target u , the random stopping time $\tau(u)$ already selects a boundary observation, with $C_{\tau(u)} = u$ and $\chi_{\tau(u)} = 0$. This is different from conditioning χ_s to be zero at a fixed magic time s . Independence of the observing driver means that, conditional on the sampled horizons, the financial process retains

its ordinary joint law at those magic times. A driver sharing financial noise would require the joint stopped law; sampling a marginal clock independently would not implement that specification. In a direct forward simulation, each vector lane must continue until its last physical target is reached. Completed lanes can be masked or compacted while unfinished lanes continue; stopping when only most lanes are complete would truncate the clock distribution.

The selected simulator uses a regular grid in mapped physical time, with spacing $1/365$ and all requested mapped market expirations inserted as exact grid nodes. Inserting an expiry splits an interval into shorter pieces. Apply the business-time and fitted physical-time maps before constructing this grid. There is no additional stock calendar grid. At every grid node the observing driver is at zero, and the financial state is retained. The positive physical increment Δu to the next node is prescribed, rather than estimated from a sampled endpoint.

The deterministic launch proxy is $y = \theta_\chi h$, so its magic launch duration is $h = y/\theta_\chi$. For $0 < \theta_\chi < 1$, choose this duration using the exact endpoint-conditioned clock mean below:

$$m(h, \theta_\chi h) = \Delta u, \quad h \sim \left[\frac{\kappa_\chi^{\theta_\chi} \Delta u}{K_{\theta_\chi}} \right]^{1/(1-\theta_\chi)}, \quad (4.261)$$

where the dimensionless small-step normalization is

$$K_{\theta_\chi} = \int_0^1 v^{-\theta_\chi} (1-v)^{-\theta_\chi} \exp\left(-\frac{\theta_\chi v}{1-v}\right) dv.$$

It can be computed numerically once per physical-driver characteristic; no additional special function or fitted coefficient is required. The positive scalar solve in (4.261) retains finite κ_χ . Precompute it for the distinct prescribed grid increments and share the result across paths. The power approximation initializes the solve; it is not substituted for the finite-mean-reversion kernel. The conditional mean is strictly increasing along $y = \theta_\chi h$, giving a unique positive root.

For comparison with the affine law, define

$$a(s) = \frac{1 - e^{-\kappa_\chi s}}{\kappa_\chi}, \quad \chi_{\tau+h}^{\text{launch}} = y = \theta_\chi h. \quad (4.262)$$

The selected scheme omits the boundary-start Gamma endpoint draw. In the underlying affine diffusion that endpoint would have shape θ_χ and scale $a(h)$, with mean $\theta_\chi a(h)$, rather than $\theta_\chi h$ at finite mean reversion. Thus the launch is a prescribed deterministic drift proxy. The omitted endpoint randomness has not been proved to be a second-order error in the grid spacing.

The exact conditional first moment used to select the proxy is

$$\begin{aligned} m(h, y) &= \mathbb{E}[C_{\tau+h} - C_\tau \mid \chi_\tau = 0, \chi_{\tau+h} = y] \\ &= \int_0^h (1 - e^{-\kappa_\chi s})^{-\theta_\chi} \left(\frac{a(h)}{a(h-s)} \right)^{\theta_\chi} \exp \left[-y \left(\frac{1}{a(h-s)} - \frac{1}{a(h)} \right) \right] ds. \end{aligned} \quad (4.263)$$

Its boundary-rate factor and density ratio include mean reversion and the physical-clock normalization of (4.243). With $A = 1 - e^{-\kappa_\chi h}$ and $r = \kappa_\chi y e^{-\kappa_\chi h} / A$, an equivalent quadrature representation is

$$m(h, y) = \frac{A^{1-\theta_\chi}}{\kappa_\chi} \int_0^1 v^{-\theta_\chi} (1-v)^{-\theta_\chi} (1-Av)^{\theta_\chi-1} \exp \left(-\frac{rv}{1-v} \right) dv. \quad (4.264)$$

The implementation instead uses Gauss–Jacobi quadrature in a scaled s/h variable with endpoint-adapted changes of variable. Small endpoints use the deficit from the zero-endpoint bridge mean; larger endpoints resolve the initial boundary. Quadrature order is explicit and stable small-argument formulas retain the finite- κ_χ kernel.

From the deterministic positive endpoint y , sample the future return to zero by

$$d = \frac{1}{\kappa_\chi} \log \left(1 + \frac{\kappa_\chi y}{G} \right), \quad G \sim \text{Gamma}(1 - \theta_\chi, 1). \quad (4.265)$$

This first-touch law is exact conditional on its starting value y . Financial evolution covers the entire launch plus return, including the positive excursion, and the update is

$$\tau \leftarrow \tau + h + d, \quad C \leftarrow C + \Delta u, \quad \chi \leftarrow 0. \quad (4.266)$$

The physical clock reaches the prescribed grid node by construction. It is a surrogate for the random boundary clock, rather than a realized boundary-clock path subsequently clipped or rescaled to an expiry. Financial states and accumulated magic durations continue across all grid nodes and expirations. The complete observing driver remains independent of financial innovations, so the financial law conditional on $h + d$ does not require a reconstructed driver bridge. At $\theta_\chi = 0$, use $h = \Delta u$, $y = d = 0$ directly, recovering the identity-clock financial branch without a clock draw.

The complete positive duration is retained as a financial event and integrated once. Compensated floating-point accumulation of magic time tracks arithmetic remainders, including events smaller than the unit in the last place of the cumulative horizon. Financial durations come from these events, rather than subtraction of two rounded horizons. This is arithmetic roundoff control; it does not alter the model's stock or clock normalization. For every path and every positive physical

grid interval, divide its complete magic duration into exactly N equal magic-time integration steps:

$$\Delta\tau_j = \frac{h+d}{N}, \quad \Delta u_j = \frac{\Delta u}{N}, \quad j = 1, \dots, N. \quad (4.267)$$

The configurable count N applies to every path, including long returns. There is no maximum magic-step length and no automatic increase of the count. Physical and magic drift and covariance contributions are combined in these same N steps; an active physical contribution does not add steps. The complete duration is integrated, with no discarded path or interpolation through a jump. The earlier automatic refinement using a maximum magic-step length was not requested and has been removed from the selected configuration.

Initial fitting uses $N = 4$ and final pricing uses $N = 16$, with clock quadrature order 32 and variance quadrature order 24 unchanged. Initial calibration uses a direct 8,192-path GPU simulation; its scores and plots are provisional. Final pricing uses a direct simulation with at least 2^{20} paths and controls disabled. Scores saved under the earlier maximum-step refinement must be distinguished from valuations at the new fixed count. The previous 256-bank configuration of 4,096 paths per bank is historical.

There are two explicit approximations: the Gamma launch endpoint is replaced by $y = \theta_\chi h$, and the random conditional boundary-clock increment is replaced by its mean to select the prescribed physical step. Their clock-law effect can persist under grid refinement. Put $p = 1 - \theta_\chi$. In the small-step limit, the return-time tail per unit physical time is

$$\frac{\kappa_\chi \theta_\chi^p}{K_{\theta_\chi} \Gamma(p+1)} (e^{\kappa_\chi x} - 1)^{-p}, \quad x > 0.$$

Comparing with (4.248) gives the dimensionless rate ratio

$$R_{\theta_\chi} = \frac{\theta_\chi^p \Gamma(\theta_\chi)}{p K_{\theta_\chi}}.$$

In general $R_{\theta_\chi} \neq 1$: the surrogate's limiting Laplace exponent is $R_{\theta_\chi} \Phi_\chi(q)$ and its limiting mean magic-to-physical rate is R_{θ_χ} . The exact-clock mean normalization remains a theoretical reference, not a proven mean-rate identity for this surrogate. No extra correction factor is applied to sampled durations. Parameter interpretation and subsequent variance tables must distinguish these two clock laws.

At fixed magic duration the reciprocal-state endpoint is a noncentral Gamma variable, so no additional rough-volatility clock or state needs to be simulated. The physical-driver return durations above are still required. Use a stock endpoint

factor with its conditional martingale normalization and the reciprocal conditional-moment construction already described above. A normalized exponential endpoint ansatz that matches the foreign first-state moment is still an approximation to the complete joint transition law; it must not be described as an exact simulation of that joint law. Retain the idiosyncratic lognormal variance and check the configurable magic-time integration count separately from the regular physical grid. Audit launch durations, sampled return tails, the exact per-path count, magic-time roundoff and the clock and variance quadratures independently. An inverse-clock benchmark additionally reports its compensation cutoff and unresolved variance. These are numerical convergence controls. They are separate from the mathematical requirement $\vartheta > 1$ for finite expected reciprocal variance; neither kind of check replaces the other.

The existing notebook is updated in place:

```
py/example_magic_only_reciprocal_forward_calibration.ipynb.
```

Its simulator is `py/magic_only_clock_forward.py`; earlier saved results remain distinguishable by their recorded clock convention and numerical configuration.

The SPX calibration retains every supplied strike, both puts and calls, and both in- and out-of-the-money options, using the established volume and bid–ask target. Zero-volume quotes may be omitted. Maturity limits are explicit: the present configuration uses seven calendar days through three calendar months. Its market snapshot contains 6,487 positive-volume quotes in 33 symbol–expiration strips. The latest request restores initial fitting at lower path counts before refinement, with a daily mapped physical grid. Present every preliminary fit with its actual sample count and control policy; final pricing and convergence checks require at least 2^{20} paths. Retain the identity-clock candidate for comparisons at common numerical settings.

Levenberg–Marquardt stages update financial coefficients and fitted physical-time interval speeds; COBYQA explores ϑ , θ_χ , and κ_χ . In accordance with the requested near-boundary policy, hold ρ fixed in LM when $|\rho| \geq 0.995$. This is an optimizer policy, not an extra financial-domain restriction. The parameter convention is `param(initial, change_scale, bounds)`: the descriptor at optimizer coordinate zero gives the initial value, and optimizer coordinates remain unbounded even for bounded model parameters. The process description precedes the editable descriptor near the top of the calibration notebook. Its triples specify all initial values, change scales and bounds; a plain number freezes that parameter. Notebook restarts default to the last saved best parameters, whose values and source are displayed separately from the editable descriptor. The descriptor’s change scales, bounds and fixed entries remain unchanged; an incompatible checkpoint raises an explicit error rather than silently resetting the fit. Later stages warm-start from

the retained candidate under the same user-specified transforms and bounds. The running fitting cell displays stage evaluation counts and budgets, the cumulative optimizer count across restarts, the current and retained errors, and every fitted expiration, refreshing its figures as the fit improves. The physical interval speeds are a nested expiration-date dictionary, generated with initial value one for every date. The earliest date is fixed at one to anchor the overall time scale; fitting all speeds together with all dimensional coefficients leaves a redundant scale direction. SPX and SPXW observations on the same date share its speed and one optimizer coordinate, while all their quotes remain in the target. Nested descriptor bounds and fixed entries apply in every stage. For a slab with magic increment h and physical increment u , combine the two channels before applying the existing reciprocal bridge. With $s = \boldsymbol{\Psi}_\zeta$, its effective state duration, stock duration, CIR shape, rate and volatility correlation are

$$\begin{aligned} D &= h + \frac{\boldsymbol{\Psi}_{\zeta,t}}{s}u, & Q &= h + \boldsymbol{\Psi}_{S,t}u, \\ a &= \frac{\vartheta h + (\boldsymbol{\Psi}_{\zeta,t} + 2\alpha_t)u/s}{D}, & b &= \frac{\beta h + \beta_t u}{D}, \\ r &= \frac{\rho\sqrt{s}h + \rho_t\sqrt{\boldsymbol{\Psi}_{\zeta,t}\boldsymbol{\Psi}_{S,t}}u}{\sqrt{sDQ}}. \end{aligned}$$

The stock/state clock ratio is Q/D . These are the same frozen-slab combinations used for the earlier independent physical/rough channels; there is no separate physical half-step or replacement clock sampler. All four physical strengths and drifts equal to zero select the previous magic-only branch exactly, including its random-draw schedule. Change scales are 0.05 for ρ and θ_χ , 0.1 for ϑ , 0.25 for physical speeds, and 0.2 times the magnitude of other initial values, using 1 for a zero initial value. Financial states, conditional payoffs and all stochastic Gamma draws, including the clock's first-return draw, use the established Torch sampler on the requested AMD/ROCm device. Only deterministic scalar clock calculations and event bookkeeping remain on CPU; redundant diagnostics share stored arrays or deterministic views. GPU availability is checked before calibration, without an automatic CPU fallback. Budget exhaustion, interruption and convergence are reported separately, with the best completed checkpoint retained.

The current workspace instructions permit explicitly labelled control variates during fitting and in provisional calibration plots, while requiring controls to be disabled in final prices, final fit plots and final convergence checks. The option-price control uses a fixed Black reference delta computed from the market ATM volatility and forward: the put coefficient is the call coefficient minus one. If X is the unadjusted normalized model stock with specified expectation 1, an option payoff Y uses $Y - b(X - 1)$ with this fixed coefficient b . Financial dynamics and

simulated stock states are unchanged. Provisional plots display both controlled and raw normalized target errors. Raw forward checks always disable controls and report the unadjusted forward discrepancy and its Monte Carlo standard error for every expiration. Final results use at least 2^{20} paths and report raw option prices, target scores, forward errors and parity errors. Do not rescale paths or forwards, shift the ATM level, or label a controlled score or plot as a final result. Every fitted expiration retains its own date, small market points, separate calls and puts, and volume panels; display ranges never select the calibration target. The latest explicit workspace instructions supersede the preceding permission to use controls in final results.

Conditional variance after calibration. The one-dimensional magic-horizon averaging formulas below apply to the special case with the four physical strengths and drifts equal to zero. With the restored physical channel, the fair area is instead $\mathbb{E}[\int \zeta^{-1}(d\tau + \boldsymbol{\Psi}_{S,t}dt)]$ under the joint physical-magic evolution. It must be evaluated or tabulated for that calibrated model; using only the terminal magic horizon would omit the physical contribution and its effect on the reciprocal state. The variance table is constructed only after selecting and saving the SPX calibration. For a fixed magic duration, let

$$\begin{aligned} G(\Delta\tau; \zeta) &= \mathbb{E}_d \left[\int_0^{\Delta\tau} \frac{dv}{\zeta_v} \middle| \zeta_0 = \zeta \right] \\ &= y \left(\frac{F_\beta(\Delta\tau)}{\zeta}; \vartheta, \boldsymbol{\Psi}_\zeta \right), \\ F_\beta(\Delta\tau) &= \frac{e^{\beta\Delta\tau} - 1}{\beta}. \end{aligned} \tag{4.268}$$

The value of F_β at zero reversion is $\Delta\tau$. The reciprocal fair-area function y is the positive mean-reverting representation in (4.123), rather than a zero-reversion formula used at the unadjusted duration. Averaging over the independent future observation clock gives

$$G_\chi(U; \zeta) = \mathbb{E}[G(K_U; \zeta)]. \tag{4.269}$$

This identity is exact for the continuous model when the expectation is finite. For the selected regular-grid surrogate, replace the exact inverse-clock horizon by the sum of its prescribed launch and sampled return durations. Averaging the analytic reciprocal fair area over that sum needs neither future stock nor future reciprocal-state paths, and retains the surrogate's complete random return distribution. It must not be identified with the exact inverse-clock result merely because the launch duration matches a conditional clock mean. Report physical grid, financial and special-function quadrature errors separately from this clock-law

difference. The inverse-clock transform and the spectral formula (4.255), when its positive-reversion assumptions and convergence checks hold, provide independent exact-clock benchmarks.

Integrating the prescribed boundary-clock distribution. For the actual future physical partition, write h_j for its prescribed launch durations and $\chi_j = \theta_\chi h_j$ for the launch endpoints. The independent subsequent return durations D_j have the same law as in the simulator:

$$D_j = \frac{1}{\kappa_\chi} \log \left(1 + \frac{\kappa_\chi \chi_j}{G_j} \right), \quad G_j \sim \text{Gamma}(1 - \theta_\chi, 1). \quad (4.270)$$

Their distribution can be expressed through the regularized upper incomplete Gamma function Q :

$$\Pr(D_j \leq d) = Q \left(1 - \theta_\chi, \frac{\kappa_\chi \chi_j}{e^{\kappa_\chi d} - 1} \right), \quad d > 0. \quad (4.271)$$

The summed horizon $K_U = \sum_j (h_j + D_j)$ is not generally Gamma. Its exact Laplace transform for this prescribed partition is

$$\mathcal{L}_U(q) = e^{-q \sum_j h_j} \prod_j \frac{\Gamma(1 - \theta_\chi + q/\kappa_\chi)}{\Gamma(1 - \theta_\chi)} \text{U} \left(\frac{q}{\kappa_\chi}, \theta_\chi, \kappa_\chi \chi_j \right). \quad (4.272)$$

Here U is Tricomi's confluent hypergeometric function, not the physical-window length. This transform is specific to the prescribed-cycle approximation; it is not the exact inverse-clock transform. The return mean is obtained by integrating the positive survival function from (4.271).

For $\beta > 0$ and $\vartheta > 1$, the existing CIR spectral representation can average the reciprocal fair area without sampling future clock or reciprocal-state paths. With $\overline{K}_U = \mathbb{E}[K_U]$ it gives

$$\begin{aligned} G_\chi(U; \zeta) &= G(\overline{K}_U; \zeta) \\ &+ \frac{2}{\psi_\zeta(\vartheta - 1)} \sum_{n=1}^{\infty} \frac{(n-1)!}{(\vartheta)_n} L_n^{(\vartheta-1)} \left(\frac{2\beta\zeta}{\psi_\zeta} \right) \left[e^{-n\beta\overline{K}_U} - \mathcal{L}_U(n\beta) \right]. \end{aligned} \quad (4.273)$$

The coefficients use generalized Laguerre polynomials and the rising factorial $(\vartheta)_n$. Evaluate the deterministic anchor with the mean-reverting reciprocal formula, and check spectral truncation, roundoff and special-function positivity. A failed check is reported, rather than replacing it with a sampled or mean-clock result. For $\theta_\chi = 0$ the clock is deterministic and the correction vanishes.

After selecting the calibrated parameters, tabulate future fair area on the observed range of ζ , including its low-state and high-state edges. Compare interpolation with independent direct evaluations, extend the table when necessary, and reject out-of-domain evaluations instead of clipping states to its edge. At regular-grid boundary observations $\chi = 0$ and the physical time is prescribed exactly; there is no carried physical-clock overshoot coordinate. The future conditional table therefore has one volatility-state coordinate, ζ , together with its configured future physical window and grid. No additional rough-volatility driver or rough-clock conditional-mean axis is introduced.

Jump variance and VIX puts. After the observation time change, the log-contract identity remains

$$-2\mathbb{E}_d \left[\log \frac{S_{T+\Delta_c}}{S_T} \middle| \zeta_T = \zeta \right] = G_\chi(U; \zeta) \quad (4.274)$$

under the discounted-stock and integrability conditions. This is the expected latent integrated variance, not generally the expected squared observed log-price jumps. An observed log jump r contributes r^2 to log-price quadratic variation, whereas the log-contract jump adjustment is $2(e^r - 1 - r)$. They agree only to second order for small jumps. A variance table and its validation must identify which of these targets it estimates; a variance-only latent simulation validates the right-hand side of (4.274).

The conditional log-contract VIX proxy is

$$\text{VIX}_T^{\log} = 100 \sqrt{\frac{G_\chi(U; \zeta_T)}{30/360}}, \quad \Delta_c = 30/365. \quad (4.275)$$

Here Δ_c specifies the thirty-calendar-day horizon in calendar-year units. The denominator $30/360$ is the separately requested annualization convention; it does not change the thirty-day window. The business clock and fitted physical-time map determine $U = u(T + \Delta_c) - u(T)$ and the actual future physical partition. For the prescribed-cycle approximation that partition is part of the clock law, so changing it is recorded separately from refining the financial integration. The accompanying comparison notebook evaluates these quantities after freezing the SPX calibration. For weekly observation conventions or a requested alternative annualization, record the actual window and denominator explicitly. Validate expected area before comparing square roots. On the same expiry-state sample, report the square root of mean benchmark variance, the square root of mean table estimate, the mean square root of the table estimate, and the market VIX future or parity-implied forward. Their differences distinguish approximation error from the model's dispersion of conditional future variance.

VIX options are reviewed after this SPX calibration and variance validation; they are not included in the current calibration target. Price puts directly from the conditional VIX states, using

$$P_T(K) = D_T \mathbb{E}[(K - \text{VIX}_T^{\log})^+]. \quad (4.276)$$

Neither the simulated VIX state nor its option prices are rescaled to match a VIX forward. Compute both market and model put implied volatilities with the same observed VIX future, or the explicitly identified parity-implied market forward, and plot them against $\log(K/F_T^{\text{market}})$. This forward is an implied-volatility reference, not a replacement for the model's own expected VIX. A model price outside the corresponding Black put bounds is reported as such, instead of changing the forward or scaling the price to make an implied volatility exist. Put payoffs emphasize the lower-VIX region, where the conditional fair-area evaluation may be more reliable; that reliability is tested rather than assumed. The comparison also does not establish exact equivalence of an ideal log contract and the finite strike portfolio used for the market VIX index.

Reciprocal dynamics in magic, physical and physical-boundary clocks

The reciprocal model in (4.258) admits an extension in which both the stock and reciprocal variance state have independent contributions from three clocks. Retain its magic-time and physical-time coefficients, and add the usual affine driver ξ in *physical* time. Its boundary clock drives an additional rough channel in both S and ζ . Setting that channel's strengths and drifts to zero recovers the existing model without changing its parameters or calibrated results. This is a new specification and an initialization for calibration, not a claim that an improved fit has already been established.

Clocks and observable states. Retain the observing driver χ_τ , normalized physical clock $t = C_\tau$, and inverse $\tau(u)$ from (4.242)–(4.243). The general specification retains the deterministic physical map $u(T)$. The completed first-week experiment fixes $u(T) = T$ in calendar time, without a holiday/business-time adjustment or fitted expiry speeds. The requested extension through the following week retains calendar time as its input, and permits the fitted physical map described below. The new driver is

$$d\xi_t = (\theta - \kappa\xi_t)dt + \sqrt{2\xi_t}dW_t^\xi, \quad 0 < \theta < 1, \quad \kappa \geq 0, \quad \xi_0 \geq 0. \quad (4.277)$$

Use the existing canonical boundary-clock normalization of (3.33) and (4.5):

$$\mathfrak{r}_t = \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^t \mathbf{1}_{\{\xi_s < \epsilon\}} ds. \quad (4.278)$$

Mean reversion changes the clock's distribution, not this local normalization. For example, at $\kappa = 0$ and $\xi_0 = 0$, $\mathbb{E}[\mathfrak{r}_t] = t^{1-\theta}/\Gamma(\theta)$; no fitted multiplier renormalizes its generated increments.

When equations are expressed in magic time, $d\mathfrak{r}$ below denotes $d(\mathfrak{r}_{C_\tau})$. During a skipped excursion of χ , $dt = 0$; therefore ξ and its rough clock do not advance. At a calendar expiry,

$$S_T = S_{\tau(u(T))}, \quad \zeta_T = \zeta_{\tau(u(T))}, \quad \xi_T = \xi_{u(T)}, \quad \chi_{\tau(u(T))} = 0. \quad (4.279)$$

Unlike χ , the observable state ξ_T need not be zero. A future valuation retains both ξ_T and ζ_T , and starts χ at zero. The separate physical-driver and rough-driver innovations are independent of each other and of the financial innovations, except for the specified stock-state correlations within each financial clock channel.

Complete drift and covariance specification. In the discounted-stock convention, the full moments are

$$\begin{aligned} \mathbb{E}_d[dS] &= 0, \\ \frac{d\langle S \rangle}{S^2} &= \frac{d\tau + \boldsymbol{\Psi}_{S,t}dt + \boldsymbol{\Psi}_S d\mathfrak{r}}{\zeta}, \\ \mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\vartheta\boldsymbol{\Psi}_\zeta - \beta\zeta\right) d\tau \\ &\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,t} + \alpha_t - \beta_t\zeta\right) dt \\ &\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,\mathfrak{r}} + \mathfrak{R} - \boldsymbol{\beth}\zeta\right) d\mathfrak{r}, \\ d\langle \zeta \rangle &= \zeta (\boldsymbol{\Psi}_\zeta d\tau + \boldsymbol{\Psi}_{\zeta,t}dt + \boldsymbol{\Psi}_{\zeta,\mathfrak{r}}d\mathfrak{r}), \\ d\langle \zeta, S \rangle &= -S(\rho\sqrt{\boldsymbol{\Psi}_\zeta} d\tau + \rho_t\sqrt{\boldsymbol{\Psi}_{\zeta,t}\boldsymbol{\Psi}_{S,t}} dt \\ &\quad + \boldsymbol{\mathfrak{N}}\sqrt{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}}\boldsymbol{\Psi}_S} d\mathfrak{r}), \\ \mathbb{E}[d\xi] &= (\theta - \kappa\xi)dt, \quad d\langle \xi \rangle = 2\xi dt, \\ \mathbb{E}[d\chi] &= (\theta_\chi - \kappa_\chi\chi)d\tau, \quad d\langle \chi \rangle = 2\chi d\tau, \\ d\langle \xi, S \rangle &= d\langle \xi, \zeta \rangle = d\langle \xi, \chi \rangle = 0, \\ d\langle \chi, S \rangle &= d\langle \chi, \zeta \rangle = 0. \end{aligned} \quad (4.280)$$

The coefficient of $d\tau/\zeta$ remains one, preserving the scale of ζ . The signs of ρ , ρ_t , and $\boldsymbol{\mathfrak{N}}$ describe variance-stock correlations; reciprocal-state-stock correlations have the opposite sign. A zero bracket between ξ and ζ does not make them

independent: ξ determines the rough clock of ζ . For $0 < \theta < 1$, $d\mathfrak{t}$ is singular relative to dt . This is consistent with a zero bracket between the physical-time driver and rough-clock innovations, while the rough innovations of S and ζ have their stated mutual covariance.

Parameter mapping. Existing parameters keep their meanings. The descriptor keys are:

Parameter	Descriptor key
ζ_0	<code>zeta_0</code>
ϑ	<code>vartheta</code>
β	<code>beta</code>
Ψ_ζ	<code>evar_zeta</code>
ρ	<code>rho</code>
θ_χ	<code>theta_chi</code>
κ_χ	<code>kappa_chi</code>
$\Psi_{S,t}$	<code>evar_S_physical</code>
$\Psi_{\zeta,t}$	<code>evar_zeta_physical</code>
α_t	<code>alpha_physical</code>
β_t	<code>beta_physical</code>
ρ_t	<code>rho_physical</code>
θ	<code>theta</code>
κ	<code>kappa</code>
ξ_0	<code>xi_0</code>
Ψ_S	<code>evar_S</code>
$\Psi_{\zeta,\mathfrak{t}}$	<code>evar_zeta_rough</code>
$\aleph$	<code>aleph</code>
$\beth$	<code>bet</code>
$\mathfrak{N}$	<code>ecorr</code>

The rough immigration is an independent excess above the compulsory one-half-diffusion floor. Thus `aleph` is not a second copy of `vartheta`, and it may remain nonzero when `evar_zeta_rough`=0. For active state diffusion the cash characteristics and stock-numeraire characteristics are

$$\begin{aligned}
\vartheta_{d,\tau} &= \vartheta, & \vartheta_{f,\tau} &= \vartheta - \frac{2\rho}{\sqrt{\Psi_\zeta}}, \\
\vartheta_{d,t} &= 1 + \frac{2\alpha_t}{\Psi_{\zeta,t}}, & \vartheta_{f,t} &= 1 + \frac{2(\alpha_t - \rho_t \sqrt{\Psi_{\zeta,t} \Psi_{S,t}})}{\Psi_{\zeta,t}}, \\
\vartheta_{d,\mathfrak{t}} &= 1 + \frac{2\aleph}{\Psi_{\zeta,\mathfrak{t}}}, & \vartheta_{f,\mathfrak{t}} &= 1 + \frac{2(\aleph - \mathfrak{N} \sqrt{\Psi_{\zeta,\mathfrak{t}} \Psi_S})}{\Psi_{\zeta,\mathfrak{t}}}.
\end{aligned} \tag{4.281}$$

They are not required to coincide. The previous shared-characteristic version in (4.136) is a constraint on a different subfamily, not a hidden restriction on these independent coefficients. Imposing it here would also tie α_t and $\mathfrak{N}$ to diffusion strengths and change the present calibrated family.

Require nonnegative diffusion strengths and stock weights, positive ζ_0 , $\vartheta > 1$, and channel correlations in $[-1, 1]$. The selected stock bridge retains its existing nonpositive leverage domain. For the new stochastic rough state channel use $\mathfrak{N} > 0$: its cash characteristic is then strictly above the inverse-moment floor one. The stock excess must satisfy $\mathfrak{N} - \mathfrak{N}\sqrt{\boldsymbol{\Psi}_{\zeta, \mathfrak{r}}\boldsymbol{\Psi}_S} \geq 0$ for its nonattainment floor. Negative leverage makes that excess no smaller than the cash excess. Retain the existing physical-channel restrictions and its strictly positive excess when both physical diffusions are active. An arbitrary additional numerical margin is not implied. If a diffusion strength vanishes, apply the corresponding drift restrictions directly, without division by that strength. In particular, at $\boldsymbol{\Psi}_{\zeta, \mathfrak{r}} = 0$, rough state diffusion and rough state–stock covariance vanish, $\mathfrak{N} \geq 0$ can provide rough drift, and $\mathfrak{M}$ is irrelevant. Both affine drivers retain their parameters under the stock numeraire because their stock brackets are zero.

Physical-clock sampling and the frozen-slab approximation. Keep the prescribed physical grid, including its explicitly configured geometric short-end refinement, and sample each existing launch-plus-return magic duration as before. Its endpoint proxy and conditional-clock-mean approximations remain those documented in (4.261)–(4.266). The new rough clock does not replace that observation-clock construction. Generate the usual affine ξ endpoints, local contact times, and canonical boundary-clock bridge increments in the same mapped physical time. Reuse the established native Torch rough-clock sampler and its mean-bridge convention; no different Gamma sampler, backend, or empirical clock normalization is introduced.

For each path and positive physical interval, retain exactly the configured number N of magic integration steps, initially four and finally sixteen. Generate ξ and its rough clock on the outer physical grid, preserving the usual rough-clock generator’s existing internal affine-step bound $\kappa\Delta t \leq 0.5$. Do not resample that clock at each financial integration step. For an outer slab, write its complete increments as $\Delta\tau, \Delta t, \Delta\mathfrak{r}$. Freeze their ratios within the slab and combine the three channels

as follows:

$$\begin{aligned}
A &= \psi_\zeta \Delta\tau + \psi_{\zeta,t} \Delta t + \psi_{\zeta,\mathfrak{r}} \Delta \mathfrak{r}, \\
W &= \Delta\tau + \psi_{S,t} \Delta t + \psi_S \Delta \mathfrak{r}, \\
I &= \frac{1}{2} \vartheta \psi_\zeta \Delta\tau + (\frac{1}{2} \psi_{\zeta,t} + \alpha_t) \Delta t + (\frac{1}{2} \psi_{\zeta,\mathfrak{r}} + \mathfrak{N}) \Delta \mathfrak{r}, \\
R &= \beta \Delta\tau + \beta_t \Delta t + \mathfrak{J} \Delta \mathfrak{r}, \\
B &= \rho \sqrt{\psi_\zeta} \Delta\tau + \rho_t \sqrt{\psi_{\zeta,t} \psi_{S,t}} \Delta t + \mathfrak{N} \sqrt{\psi_{\zeta,\mathfrak{r}} \psi_S} \Delta \mathfrak{r}.
\end{aligned} \tag{4.282}$$

Use the reference state strength $\psi_\zeta > 0$ and define

$$D = A/\psi_\zeta, \quad \vartheta_{\text{step}} = 2I/A, \quad \beta_{\text{step}} = R/D, \quad \rho_{\text{step}} = B/\sqrt{AW}. \tag{4.283}$$

The positive magic contribution ensures $A, W > 0$ on these substeps. Cauchy-Schwarz gives $|\rho_{\text{step}}| \leq 1$. The exact noncentral Gamma transition for this frozen CIR slab advances ζ over duration D , with state diffusion strength ψ_ζ , cash characteristic ϑ_{step} , and integrated reversion R . Its stock-numeraire characteristic is $(2I - 2B)/A$.

Integrate the frozen slab in exactly N equal financial steps of state duration D/N and stock duration W/N . Equivalently, each such step receives $\Delta\tau/N$, $\Delta t/N$, and $\Delta \mathfrak{r}/N$ with the same frozen coefficients. This resolves the financial transition without claiming to resolve the ordering of the rough clock inside the outer slab.

Apply the existing reciprocal conditional-stock construction with stock clock duration W . Its fair inverse-state area is the state-clock integral multiplied by W/D . Retain the analytic conditional martingale normalizer, the foreign first-state-moment condition, and the idiosyncratic lognormal residual variance; none is replaced by a sample-mean correction. The step still approximates the unresolved ordering of the three clocks. The varying characteristic in (4.283) is an explicit part of the existing frozen-clock approximation, not a claim that a time-varying three-clock CIR transition has a single exact global shape. The mean-bridge rough clock also retains its established higher-moment approximation. Physical-grid, magic-step, and variance quadrature checks remain separate.

Embedding, scaling, and reported results. When $\psi_S = \psi_{\zeta,\mathfrak{r}} = \mathfrak{N} = \mathfrak{J} = 0$, bypass the new driver draws and delegate directly to the previous simulator. This retains its random inputs and prices, rather than only equality in law. At $\theta_\chi = 0$, use $C_\tau = \tau$ as before: the magic and physical financial channels coalesce, the observation jumps disappear, and the remaining physical-boundary clock still supplies rough stock and state motion. Its boundary scaling is $\mathfrak{r}_t \sim t^{1-\theta}$, with Brownian marginal exponent $(1 - \theta)/2$; positive initial state and mean reversion introduce contact delays and crossovers.

This construction differs from the historical three-clock model below, whose ξ evolves in magic time. In particular, that model's ratio $(1 - \theta)/(1 - \theta_\chi)$ is not the

scaling power of the new physical-time rough driver. In engineering time units, $[\xi] = [\zeta] = T$, $[\mathfrak{r}] = T^{1-\theta}$, $[\kappa] = T^{-1}$, and $[\boldsymbol{\psi}_S] = [\boldsymbol{\psi}_{\zeta, \mathfrak{r}}] = [\boldsymbol{\mathfrak{N}}] = T^\theta$, $[\boldsymbol{\beth}] = T^{\theta-1}$; the driver characteristics and correlations are dimensionless. Normalized engineering units remain distinct from the short-horizon physical/magic scaling controlled by θ_χ .

Calendar-based expiry speeds. For calendar expiries $0 = T_0 < T_1 < \dots < T_m$, an optional positive piecewise-linear map assigns one speed r_i to each preceding interval:

$$u(T) = \sum_{j < i} r_j(T_j - T_{j-1}) + r_i(T - T_{i-1}), \quad T_{i-1} \leq T \leq T_i, \quad r_i > 0, \quad r_1 = 1. \quad (4.284)$$

There is no additional holiday/business-time conversion in this test. All initial speeds are one, so the initial map is the identity. The first fixed speed anchors the overall time scale. Keep the editable speeds in a nested `time_speed` dictionary keyed by expiration date: its first entry is the number one, and each subsequent entry is `param(1, change_scale, [0])`. The same date has one speed, independent of option right. This map supplies the physical grid both to the new ξ and rough clock and to the existing physical channel and prescribed magic-clock sampler. It is not a rescaling of sampled magic durations. Market quotes, forwards, and the calendar maturity used to report implied volatility are not rescaled by this map.

The separate implementation and calibration notebook retain the old version and its checkpoints. The completed first-week result also remains separately recorded. The wider-expiry experiment starts from its saved raw-final financial and driver parameters, with the added expiry speeds initialized to one. Reprice this candidate on the whole new quote set before optimization; a saved score from four dates is not an error for the extended target. Initial low-path, labelled-control trials do not certify final results. Final scores and plots require at least 2^{20} paths, controls disabled, and per-expiry raw forward and parity errors. The completed initial test uses Monday 5 October 2026 EOD data and the four daily SPXW expirations 6–9 October, with the identity calendar-time map. Its raw normalized error was 1.14604170 at 2^{20} paths and sixteen financial steps, with controls disabled. The retained candidate has zero added rough strengths and drifts, so this result lies in the magic/physical subfamily; the local search does not establish a global optimum.

The requested wider test retains the same snapshot and adds the five daily expirations 12–16 October, for nine SPXW dates in total, with the optional fitted physical map in (4.284). Both tests use reported EOD close premiums as targets and retain source bid–ask half-spreads and positive-volume weights, both rights and all supplied strikes. The datafeed, objective form, editable descriptor conventions,

and numerical backend remain the existing ones. No extended-scope calibration result is asserted here; its score and every expiry's fitted calendar-to-model time must be reported after repricing.

Reciprocal dynamics with a magic-time rough driver

A separate variant moves the affine rough-clock driver from physical into magic time. Both the stock and reciprocal variance state retain magic, physical and rough financial contributions. No additional Heston stock-variance term or driver-stock covariance is included in this selected variant. The physical-driver model in (4.280), its implementation and its calibrated results remain intact.

Clocks and observable states. Retain the observation driver χ_τ , its normalized physical boundary clock $t = C_\tau$, and its inverse from (4.242)–(4.243). The added driver now evolves in the same magic time:

$$d\xi_\tau = (\theta - \kappa\xi_\tau)d\tau + \sqrt{2\xi_\tau}dW_\tau^\xi, \quad 0 < \theta < 1, \quad \kappa \geq 0, \quad \xi_0 \geq 0. \quad (4.285)$$

Use the same canonical local-time normalization, with magic time as its integration variable:

$$\mathfrak{r}_\tau = \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^\tau \mathbf{1}_{\{\xi_s < \epsilon\}} ds. \quad (4.286)$$

At $\kappa = 0$, $\xi_0 = 0$, $\mathbb{E}[\mathfrak{r}_\tau] = \tau^{1-\theta}/\Gamma(\theta)$. Mean reversion changes the distribution, not this normalization. The physical observation driver has $\chi_0 = 0$, with $0 < \theta_\chi < 1$ and $\kappa_\chi > 0$ for its active boundary clock; at $\theta_\chi = 0$ retain the direct convention $C_\tau = \tau$.

For the deterministic physical target $u(T)$, the observable states are

$$S_T = S_{\tau(u(T))}, \quad \zeta_T = \zeta_{\tau(u(T))}, \quad \xi_T = \xi_{\tau(u(T))}, \quad \chi_{\tau(u(T))} = 0. \quad (4.287)$$

Thus the rough-driver state is retained between observations. Unlike the physical-driver version, ξ and its boundary clock continue through skipped magic excursions of χ , even when $dt = 0$. Both drivers have independent innovations, and their innovations are independent of the financial innovations in this variant. Dependence through the clocks remains: zero cross-brackets do not make the observed financial state independent of its clock drivers.

Complete drift and covariance specification. In the discounted-stock convention, the full moments are

$$\begin{aligned}
\mathbb{E}_d[dS] &= 0, \\
\frac{d\langle S \rangle}{S^2} &= \frac{d\tau + \boldsymbol{\Psi}_{S,t}dt + \boldsymbol{\Psi}_S d\mathfrak{r}}{\zeta}, \\
\mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\partial\boldsymbol{\Psi}_\zeta - \beta\zeta \right) d\tau \\
&\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,t} + \alpha_t - \beta_t\zeta \right) dt \\
&\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,\mathfrak{r}} + \mathfrak{R} - \mathfrak{J}\zeta \right) d\mathfrak{r}, \\
d\langle \zeta \rangle &= \zeta (\boldsymbol{\Psi}_\zeta d\tau + \boldsymbol{\Psi}_{\zeta,t}dt + \boldsymbol{\Psi}_{\zeta,\mathfrak{r}}d\mathfrak{r}), \\
d\langle \zeta, S \rangle &= -S(\rho\sqrt{\boldsymbol{\Psi}_\zeta} d\tau + \rho_t\sqrt{\boldsymbol{\Psi}_{\zeta,t}\boldsymbol{\Psi}_{S,t}} dt \\
&\quad + \mathfrak{N}\sqrt{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}}\boldsymbol{\Psi}_S} d\mathfrak{r}), \\
\mathbb{E}[d\xi] &= (\theta - \kappa\xi)d\tau, \quad d\langle \xi \rangle = 2\xi d\tau, \\
\mathbb{E}[d\chi] &= (\theta_\chi - \kappa_\chi\chi)d\tau, \quad d\langle \chi \rangle = 2\chi d\tau, \\
d\langle \xi, S \rangle &= d\langle \xi, \zeta \rangle = d\langle \xi, \chi \rangle = 0, \\
d\langle \chi, S \rangle &= d\langle \chi, \zeta \rangle = 0.
\end{aligned} \tag{4.288}$$

The coefficient of $d\tau/\zeta$ remains one, preserving the scale of ζ . The signs of ρ , ρ_t , and $\mathfrak{N}$ describe variance–stock correlations; reciprocal-state–stock correlations have the opposite sign. A zero bracket between ξ and ζ does not make them independent: ξ determines the rough clock of ζ . Here $d\mathfrak{r}$ means the boundary clock of ξ in magic time. The driver–stock zero brackets are specified restrictions of this variant, not consequences of clock singularity. The independent financial innovations retain the stated stock–state correlations.

Parameter mapping. Existing parameters keep their meanings. The descriptor keys are:

Parameter	Descriptor key
ζ_0	<code>zeta_0</code>
ϑ	<code>vartheta</code>
β	<code>beta</code>
Ψ_ζ	<code>evar_zeta</code>
ρ	<code>rho</code>
θ_χ	<code>theta_chi</code>
κ_χ	<code>kappa_chi</code>
$\Psi_{S,t}$	<code>evar_S_physical</code>
$\Psi_{\zeta,t}$	<code>evar_zeta_physical</code>
α_t	<code>alpha_physical</code>
β_t	<code>beta_physical</code>
ρ_t	<code>rho_physical</code>
θ	<code>theta</code>
κ	<code>kappa</code>
ξ_0	<code>xi_0</code>
Ψ_S	<code>evar_S</code>
$\Psi_{\zeta,\mathfrak{r}}$	<code>evar_zeta_rough</code>
$\aleph$	<code>aleph</code>
$\beth$	<code>bet</code>
$\mathfrak{N}$	<code>ecorr</code>

The rough immigration is an independent excess above the compulsory one-half-diffusion floor. Thus `aleph` is not a second copy of `vartheta`, and it may remain nonzero when `evar_zeta_rough`=0. For active state diffusion the cash characteristics and stock-numeraire characteristics are

$$\begin{aligned}
\vartheta_{d,\tau} &= \vartheta, & \vartheta_{f,\tau} &= \vartheta - \frac{2\rho}{\sqrt{\Psi_\zeta}}, \\
\vartheta_{d,t} &= 1 + \frac{2\alpha_t}{\Psi_{\zeta,t}}, & \vartheta_{f,t} &= 1 + \frac{2(\alpha_t - \rho_t \sqrt{\Psi_{\zeta,t} \Psi_{S,t}})}{\Psi_{\zeta,t}}, \\
\vartheta_{d,\mathfrak{r}} &= 1 + \frac{2\aleph}{\Psi_{\zeta,\mathfrak{r}}}, & \vartheta_{f,\mathfrak{r}} &= 1 + \frac{2(\aleph - \mathfrak{N} \sqrt{\Psi_{\zeta,\mathfrak{r}} \Psi_S})}{\Psi_{\zeta,\mathfrak{r}}}.
\end{aligned} \tag{4.289}$$

They are not required to coincide. The previous shared-characteristic version in (4.136) is a constraint on a different subfamily, not a hidden restriction on these independent coefficients. Imposing it here would also tie α_t and $\aleph$ to diffusion strengths and change the present calibrated family.

Require nonnegative diffusion strengths and stock weights, positive ζ_0 , $\vartheta > 1$, and channel correlations in $[-1, 1]$. The selected stock bridge retains its existing nonpositive leverage domain. For the new stochastic rough state channel use $\aleph > 0$:

its cash characteristic is then strictly above the inverse-moment floor one. The stock excess must satisfy $\mathfrak{N} - \mathfrak{N}\sqrt{\boldsymbol{\Psi}_{\zeta, \mathfrak{r}}\boldsymbol{\Psi}_S} \geq 0$ for its nonattainment floor. Negative leverage makes that excess no smaller than the cash excess. Retain the existing physical-channel restrictions and its strictly positive excess when both physical diffusions are active. An arbitrary additional numerical margin is not implied. If a diffusion strength vanishes, apply the corresponding drift restrictions directly, without division by that strength. In particular, at $\boldsymbol{\Psi}_{\zeta, \mathfrak{r}} = 0$, rough state diffusion and rough state–stock covariance vanish, $\mathfrak{N} \geq 0$ can provide rough drift, and $\mathfrak{N}$ is irrelevant. Both affine drivers retain their parameters under the stock numeraire because their stock brackets are zero.

Prescribed physical grid and magic-time rough sampling. Retain the existing prescribed physical grid, including its editable geometric short-end refinement and every expiry knot. Generate the observation clock using the same prescribed endpoint, conditional finite-mean-reversion clock mean and random Gamma return construction. The complete launch-plus-return duration is retained for every path. Its existing approximation is not replaced by exact inverse-clock sampling or a new launch distribution.

Generate the added rough state over each path’s complete retained magic duration. Use the stored positive cycle duration itself as $\Delta\tau$, consistently with the financial kernel. A difference of cumulative magic horizons is only a consistency check: accumulation roundoff can conceal an increment smaller than the horizon’s unit of precision. Neither the clock increment nor its normalization is rescaled to correct that accumulation.

Reuse the existing native Torch affine rough-clock step, carrying each lane’s endpoint ξ and accumulated $\mathfrak{r}$ to the next outer slab. To preserve its existing clock-only affine bound, use the vector-wide subdivision count

$$n = \max \left(1, \left\lceil \frac{\kappa \max_{\text{paths}}(\Delta\tau)}{0.5} \right\rceil \right), \quad (4.290)$$

and give each lane n calls of duration $\Delta\tau/n$. The longest magic lane controls this rough-driver subdivision and the count is reported. This is separate from the exactly configured number N of financial integration steps; no cap or automatic financial-step refinement is introduced. The driver endpoints and rough increments are generated over the outer slab, rather than resampled at its financial substeps.

For the complete increments $\Delta\tau$, Δt and $\Delta\mathfrak{r}$, freeze their ratios and combine

the same three financial channels:

$$\begin{aligned}
A &= \psi_\zeta \Delta\tau + \psi_{\zeta,t} \Delta t + \psi_{\zeta,\mathfrak{r}} \Delta \mathfrak{r}, \\
W &= \Delta\tau + \psi_{S,t} \Delta t + \psi_S \Delta \mathfrak{r}, \\
I &= \frac{1}{2} \vartheta \psi_\zeta \Delta\tau + (\frac{1}{2} \psi_{\zeta,t} + \alpha_t) \Delta t + (\frac{1}{2} \psi_{\zeta,\mathfrak{r}} + \aleph) \Delta \mathfrak{r}, \\
R &= \beta \Delta\tau + \beta_t \Delta t + \beth \Delta \mathfrak{r}, \\
B &= \rho \sqrt{\psi_\zeta} \Delta\tau + \rho_t \sqrt{\psi_{\zeta,t} \psi_{S,t}} \Delta t + \mathfrak{N} \sqrt{\psi_{\zeta,\mathfrak{r}} \psi_S} \Delta \mathfrak{r}.
\end{aligned} \tag{4.291}$$

Use the reference state strength $\psi_\zeta > 0$ and define

$$D = A/\psi_\zeta, \quad \vartheta_{\text{step}} = 2I/A, \quad \beta_{\text{step}} = R/D, \quad \rho_{\text{step}} = B/\sqrt{AW}. \tag{4.292}$$

The positive magic contribution ensures $A, W > 0$ on these substeps. Cauchy-Schwarz gives $|\rho_{\text{step}}| \leq 1$. The exact noncentral Gamma transition for this frozen CIR slab advances ζ over duration D , with state diffusion strength ψ_ζ , cash characteristic ϑ_{step} , and integrated reversion R . Its stock-numeraire characteristic is $(2I - 2B)/A$.

Integrate the frozen slab in exactly N equal financial steps of state duration D/N and stock duration W/N . Equivalently, each such step receives $\Delta\tau/N$, $\Delta t/N$, and $\Delta \mathfrak{r}/N$ with the same frozen coefficients. This resolves the financial transition without claiming to resolve the ordering of the rough clock inside the outer slab.

Apply the existing reciprocal conditional-stock construction with stock clock duration W . Its fair inverse-state area is the state-clock integral multiplied by W/D . Retain the analytic conditional martingale normalizer, the foreign first-state-moment condition, and the idiosyncratic lognormal residual variance; none is replaced by a sample-mean correction. The step still approximates the unresolved ordering of the three clocks. The varying characteristic in (4.292) is an explicit part of the existing frozen-clock approximation, not a claim that a time-varying three-clock CIR transition has a single exact global shape. The mean-bridge rough clock also retains its established higher-moment approximation. Physical-grid, magic-step, and variance quadrature checks remain separate.

Embedding and time dimensions. When $\psi_S = \psi_{\zeta,\mathfrak{r}} = \aleph = \beth = 0$, bypass the added rough draws and delegate to the existing magic/physical simulator. This preserves the previous random inputs and prices exactly. At $\theta_\chi = 0$, magic and physical time coincide; the observation jumps disappear, while the remaining magic-boundary clock still supplies its rough stock and state contributions.

The rough-driver units are magic-time units:

$$\begin{aligned}
[\xi] &= [\chi] = [\zeta] = [\tau], & [\kappa] &= [\kappa_\chi] = [\beta] = [\tau]^{-1}, \\
[\mathfrak{r}] &= [\tau]^{1-\theta}, & [\boldsymbol{\Psi}_S] &= [\boldsymbol{\Psi}_{\zeta, \mathfrak{r}}] = [\mathfrak{N}] = [\tau]^\theta, \\
[\mathfrak{z}] &= [\tau]^{\theta-1}, & [\boldsymbol{\Psi}_{S,t}] &= [\boldsymbol{\Psi}_{\zeta,t}] = [\alpha_t] = [\tau]/[t], \\
[\beta_t] &= [t]^{-1}.
\end{aligned} \tag{4.293}$$

$\boldsymbol{\Psi}_\zeta$, the Feller characteristics and correlations are dimensionless. Engineering normalization of the physical clock does not identify its boundary scaling with magic-time scaling. In the zero-state scaling comparison, $\mathfrak{r} \sim \tau^{1-\theta}$ and $t \sim \tau^{1-\theta_\chi}$ give the nominal rough-clock power relative to physical time

$$\frac{1 - \theta}{1 - \theta_\chi}. \tag{4.294}$$

The corresponding clocked-Brownian power is half this ratio. Positive initial state, mean reversion and the deterministic fitted time map introduce crossover scales; the observed process has jumps, so this comparison does not assign it a classical local Hurst exponent.

Separate full-month calibration experiment. Retain the positive, calendar-based physical map (4.284), with one date-keyed interval speed, unit descriptor defaults and the first fixed unit anchor. There is no new business/holiday adjustment. The map supplies the physical targets to the observation-clock sampler; those generated magic durations then advance the new rough driver. External market maturities, carry and implied-volatility annualization remain unchanged.

Use Monday 5 October 2026 EOD SPXW data at all twenty-three expirations 6–9, 12–16, 19–23 and 26–30 October, and 2–5 November. Preserve all 4,863 positive-volume quotes, both rights and ITM/OTM strikes, the reported close premium target, observed bid–ask half-spreads and source volume weights. Zero-volume quotes have zero target weight; no new strike cutoff is applied.

The new notebook visibly initializes only the twelve rough-driver independent base parameters and all twenty-three dated speeds from the previous full-month confirmed vector. It imports none of that variant’s prices, error, optimizer history or checkpoint. The eight rough parameters have fresh editable initial values: $\theta = 0.8$, $\kappa = 10$, $\xi_0 = 0.002$, $\boldsymbol{\Psi}_S = 0.02$, $\boldsymbol{\Psi}_{\zeta, \mathfrak{r}} = 35.4145762448321$, $\mathfrak{N} = 13.441311872884146$, $\mathfrak{z} = 0.14107684757865283$ and $\mathfrak{N} = -0.95$. Existing descriptor scales and bounds remain in use, including the approved κ change scale 100.

Retain the editable fitting ladder 32,768, 131,072 and 1,048,576 paths, with exactly four, sixteen and sixteen financial steps, respectively. At each rung run

the existing financial/time LM and rough/shape/time COBYQA blocks with their configured 160-evaluation budgets. Include all twenty-two free speeds in both applicable global blocks. The approved physical correlation is fitted in COBYQA, with its separate -0.05 starting trial only at the first rung. Later rungs retain the attained correlation. Lower-path controls and plots remain explicitly provisional; the final million-path fitting rung and independent final pricing use seed 202611001 and controls disabled.

The separate implementation and notebook have their own result directory and resume only their own saved best. Keep the model specification before the editable descriptor, live current/retained scores and counts, every fitted expiry plot, the physical-time map and raw forward diagnostics. The existing AMD/ROCm Torch environment and shared Jupyter server are retained. No calibration improvement or convergence result for this new variant is asserted here.

Two affine stock contributions in the magic-time rough-driver model

This separate extension adds stock contributions from both affine clock drivers to (4.288). It preserves all three reciprocal channels and their independent coefficients. The earlier magic-time rough-driver model, its implementation and calibration remain the nested comparison. Use the normalization of (4.57): the new strengths are $\mathcal{H}_\xi, \mathcal{H}_\chi \geq 0$, and their independent component correlations are $\rho_\xi, \rho_\chi \in [-1, 1]$.

Full cash-measure specification. The complete discounted-stock moments are

$$\begin{aligned}
\mathbb{E}_d[dS] &= 0, \\
\frac{d\langle S \rangle}{S^2} &= \frac{d\tau + \boldsymbol{\psi}_{S,t}dt + \boldsymbol{\psi}_S d\mathfrak{r}}{\zeta} + \frac{\mathcal{H}_\xi \xi + \mathcal{H}_\chi \chi}{2} d\tau, \\
\mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\vartheta \boldsymbol{\psi}_\zeta - \beta \zeta\right) d\tau + \left(\frac{1}{2}\boldsymbol{\psi}_{\zeta,t} + \alpha_t - \beta_t \zeta\right) dt \\
&\quad + \left(\frac{1}{2}\boldsymbol{\psi}_{\zeta,\mathfrak{r}} + \mathfrak{N} - \mathfrak{J} \zeta\right) d\mathfrak{r}, \\
d\langle \zeta \rangle &= \zeta (\boldsymbol{\psi}_\zeta d\tau + \boldsymbol{\psi}_{\zeta,t}dt + \boldsymbol{\psi}_{\zeta,\mathfrak{r}}d\mathfrak{r}), \\
d\langle \zeta, S \rangle &= -S(\rho \sqrt{\boldsymbol{\psi}_\zeta} d\tau + \rho_t \sqrt{\boldsymbol{\psi}_{\zeta,t} \boldsymbol{\psi}_{S,t}} dt \\
&\quad + \mathfrak{N} \sqrt{\boldsymbol{\psi}_{\zeta,\mathfrak{r}} \boldsymbol{\psi}_S} d\mathfrak{r}), \\
\mathbb{E}_d[d\xi] &= (\theta - \kappa \xi) d\tau, \quad d\langle \xi \rangle = 2\xi d\tau, \\
\mathbb{E}_d[d\chi] &= (\theta_\chi - \kappa_\chi \chi) d\tau, \quad d\langle \chi \rangle = 2\chi d\tau, \\
d\langle \xi, S \rangle &= \rho_\xi S \sqrt{\mathcal{H}_\xi} \xi d\tau, \\
d\langle \chi, S \rangle &= \rho_\chi S \sqrt{\mathcal{H}_\chi} \chi d\tau, \\
d\langle \xi, \chi \rangle &= d\langle \xi, \zeta \rangle = d\langle \chi, \zeta \rangle = 0.
\end{aligned} \tag{4.295}$$

The affine drivers retain independent Brownian innovations. Each new stock contribution shares only its own driver's innovation and an independent residual innovation. In particular,

$$\frac{dS}{S} \Big|_{\mathcal{H}_j} = \sqrt{\frac{\mathcal{H}_{jj}}{2}} \left(\rho_j dW^j + \sqrt{1 - \rho_j^2} dW^{j,\perp} \right), \quad j \in \{\xi, \chi\}. \tag{4.296}$$

Consequently the two new ρ_j are Brownian correlations with their respective affine stock components. They are distinct from the three existing reciprocal variance-stock correlations. Their covariances in (4.295) contain no additional factor $1/2$ or $\sqrt{2}$. Independent component choices in $[-1, 1]$ give a positive semidefinite full covariance; a constraint on the sum of their squared correlations is unnecessary. Their correlations with the total stock innovation are diluted by the other stock-variance contributions.

Clocks, observations and parameter mapping. The canonical rough clock remains the boundary clock of ξ in magic time, and the physical clock remains

the normalized boundary clock of χ :

$$\begin{aligned}\mathfrak{r}_\tau &= \lim_{\epsilon \downarrow 0} \frac{\theta(1-\theta)}{\epsilon^\theta} \int_0^\tau \mathbf{1}_{\{\xi_s < \epsilon\}} ds, \\ \mathfrak{r}_\tau^\chi &= \lim_{\epsilon \downarrow 0} \frac{\theta_\chi(1-\theta_\chi)}{\epsilon^{\theta_\chi}} \int_0^\tau \mathbf{1}_{\{\chi_s < \epsilon\}} ds, \\ C_\tau &= a_\chi \mathfrak{r}_\tau^\chi, \quad a_\chi = \frac{\Gamma(\theta_\chi)}{(1-\theta_\chi)\kappa_\chi^{\theta_\chi}}, \quad \tau(u) = \inf\{s : C_s > u\}.\end{aligned}\tag{4.297}$$

Here a_χ is fixed by the cash specification. Use $0 < \theta < 1$, $0 < \theta_\chi < 1$ for active boundary clocks; the existing normalized physical-clock construction uses $\kappa_\chi > 0$. Start at $\chi_0 = 0$ at each observed valuation date and retain ξ_0 . At $\theta_\chi = 0$, set $C_\tau = \tau$ and $\chi = 0$ directly, so the $\mathcal{H}_\chi$ contribution vanishes.

For the configured deterministic physical target $u(T)$,

$$S_T = S_{\tau(u(T))}, \quad \zeta_T = \zeta_{\tau(u(T))}, \quad \xi_T = \xi_{\tau(u(T))}, \quad \chi_{\tau(u(T))} = 0. \tag{4.298}$$

The new χ stock contribution generally has positive integrated variance during a skipped excursion, despite its zero observed endpoints. The rough driver and its clock continue over the entire retained magic duration, including the χ return.

All twenty descriptor mappings in the preceding variant remain unchanged. The four added entries are

Symbol	Descriptor key	Dimension
$\mathcal{H}_\xi$	H_xi	$[\tau]^{-2}$
$\mathcal{H}_\chi$	H_chi	$[\tau]^{-2}$
ρ_ξ	rho_xi	dimensionless
ρ_χ	rho_chi	dimensionless

These strengths are distinct from a Hurst exponent. The engineering normalization of physical time does not change their intrinsic magic-time dimensions or the boundary scaling discussed after (4.293).

Stock-numeraire change and the fixed physical clock. The cross-brackets give separate affine drift shifts:

$$\begin{aligned}\mathbb{E}_f[d\xi] &= [\theta - \kappa^f \xi] d\tau, \quad \kappa^f = \kappa - \rho_\xi \sqrt{\mathcal{H}_\xi}, \\ \mathbb{E}_f[d\chi] &= [\theta_\chi - \kappa_\chi^f \chi] d\tau, \quad \kappa_\chi^f = \kappa_\chi - \rho_\chi \sqrt{\mathcal{H}_\chi}.\end{aligned}\tag{4.299}$$

The two affine Feller characteristics θ, θ_χ do not change. The reciprocal stock drift remains

$$\begin{aligned}\mathbb{E}_f[d\zeta] &= \left(\frac{1}{2}\vartheta\boldsymbol{\Psi}_\zeta - \rho\sqrt{\boldsymbol{\Psi}_\zeta} - \beta\zeta\right) d\tau \\ &\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,t} + \alpha_t - \rho_t\sqrt{\boldsymbol{\Psi}_{\zeta,t}\boldsymbol{\Psi}_{S,t}} - \beta_t\zeta\right) dt \\ &\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,\mathfrak{t}} + \boldsymbol{\aleph} - \boldsymbol{\mathfrak{n}}\sqrt{\boldsymbol{\Psi}_{\zeta,\mathfrak{t}}\boldsymbol{\Psi}_S} - \boldsymbol{\beth}\zeta\right) d\mathfrak{t}.\end{aligned}\tag{4.300}$$

Thus the three reciprocal characteristics retain (4.289). The new affine strengths do not shift them because their Brownian stock contributions have zero bracket with ζ .

The path functional $C_\tau = a_\chi \mathfrak{R}_\tau^\chi$ and target $u(T)$ are the same under both numeraire: do not replace the cash κ_χ in a_χ by κ_χ^f . Doing so would change the physical observation rule. The law of the stopped magic duration changes under the stock numeraire because χ is now correlated with stock. Accordingly, the earlier argument based on an independent random horizon cannot be used to prove the stopped stock martingale. Its conditional construction must retain this change of duration law, as well as suitable integrability.

Retained first-return information. The existing prescribed physical interval selects a deterministic launch duration h and endpoint $y = \theta_\chi h$ from the cash conditional physical-clock mean. Retain its random return duration D to zero. For $0 < \theta_\chi < 1$ and $\kappa_\chi \geq 0$, its density is

$$f_{\kappa_\chi}(D | y) = \frac{\kappa_\chi(\kappa_\chi y)^{1-\theta_\chi} \exp\{\kappa_\chi D - \kappa_\chi y/(e^{\kappa_\chi D} - 1)\}}{\Gamma(1 - \theta_\chi)(e^{\kappa_\chi D} - 1)^{2-\theta_\chi}}. \quad (4.301)$$

Its zero-rate limit is $y^{1-\theta_\chi} e^{-y/D} / [\Gamma(1 - \theta_\chi) D^{2-\theta_\chi}]$. The retained duration stock factor must include

$$M_{\chi,D} = \frac{f_{\kappa_\chi^f}(D | y)}{f_{\kappa_\chi}(D | y)}. \quad (4.302)$$

Both densities integrate to one in the existing nonnegative-rate first-return helper domain; hence the factor is analytically normalized over the actual sampled cash return. Final $\chi = 0$ alone contains neither this likelihood nor the excursion's integrated area.

Reuse the existing conditional first-return affine transitions to retain χ inside the return, and integrate their positive conditional mean areas at the existing quadrature nodes. Exactly N financial substeps partition each complete launch-plus-return cycle. Its prescribed launch uses $\chi(s) = \theta_\chi s$, with area $\theta_\chi h^2/2$. The launch stock variance is therefore $\mathcal{H}_\chi \theta_\chi h^2/4$; the retained launch approximation assigns it independent compensated Gaussian noise. It discards launch leverage. An unconditional CIR launch endpoint likelihood must not be applied to this deterministic endpoint: it is not sampled from that CIR marginal and would not be normalized under the selected launch law.

For each return substep, reuse the analytically normalized noncentral Gamma endpoint factor, with the foreign conditional first-return mean computed at κ_χ^f . The remaining Gaussian variance is $(1 - \rho_\chi^2) \mathcal{H}_\chi/2$ times the conditional return area. The complete χ factor retains both these bridge factors and (4.302); it does not set the conditional return-duration factor to one.

Rough-driver endpoint approximation and exact nesting. Retain the existing ξ rough-clock bank, its sampled endpoints and complete outer magic durations. The separate ξ stock factor uses the existing normalized affine endpoint factor with cash shape θ and the foreign mean determined by κ^f . Its positive forward-mean area is

$$\mathbb{E}_d \left[\int_0^L \xi_s ds \mid \xi_0 \right] = \xi_0 \frac{1 - e^{-\kappa L}}{\kappa} + \frac{\theta}{\kappa} \left(L - \frac{1 - e^{-\kappa L}}{\kappa} \right), \quad (4.303)$$

with limit $\xi_0 L + \theta L^2/2$ at $\kappa = 0$. Its residual variance is $(1 - \rho_\xi^2) \mathcal{H}_\xi/2$ times this area. This uses retained outer endpoints; it does not introduce a new ξ path sampler or claim to recover all continuous joint rough-clock, area and return fluctuations.

The existing reciprocal factor stays conditionally normalized given its retained clock inputs. The new endpoint and duration factors multiply it, while independent compensated Gaussian variances add. These are analytic discrete first-moment normalizers, not sample forward corrections. They do not make the complete finite-grid joint law exact. At $\mathcal{H}_\xi = \mathcal{H}_\chi = 0$, delegate directly to the old simulator, including its random inputs, snapshots and outputs; both new correlations are then irrelevant. There is no new state or clock.

Mathematical and helper domains; implementation status. Retain the old reciprocal positivity and finite-inverse-moment conditions in both numeraires. Add only nonnegative $\mathcal{H}_j$ and $|\rho_j| \leq 1$ to the covariance specification. The existing numerical first-return helpers require $\kappa_\chi, \kappa_\chi^f \geq 0$ and strictly positive return starts and durations. The native ξ sampler retains its nonnegative cash-rate domain. Its endpoint factor requires a finite second moment of the selected exponential tilt; failure is reported rather than clipped. These are specific helper domains, not a new arbitrary parameter margin or a claim that every negative affine reversion rate is mathematically inadmissible. The active physical-clock normalization retains its positive cash rate. The initial separate implementation supports $\rho_\xi, \rho_\chi \in [-1, 0]$, retaining the selected negative-leverage domain; this narrower support does not make positive component correlation mathematically inadmissible.

The separate core, market and calibration adapters and notebook are implemented under model identifier `magic_xi_heston_reciprocal_v1`; its market adapter appends `_prescribed_physical_cycles_v1`. Thirteen core and six adapter CPU tests pass. A labelled GPU diagnostic on the shared AMD environment used 1,024 test paths at four and sixteen financial substeps. It verified zero-strength bit-wise nesting, finite outputs with both new channels active, GPU stock tensors and the exact configured χ step counts. The separate calibration notebooks and saved results record current errors, parameter vectors, evaluation counts and numerical settings. Raw confirmations use at least 2^{20} paths, sixteen financial substeps and

controls disabled; provisional controlled fitting remains separately labelled. Implementation tests and completed fitting passes do not establish convergence or optimality. The old retained fit remains the exact zero-strength comparison, with raw forward, parity, foreign-state moments and grid/substep checks reported separately.

Reciprocal dynamics with the joint boundary clock of two magic-time drivers

A third separate variant retains both affine drivers in magic time, but replaces the rough clock of ξ by the canonical boundary clock of $Y = \chi + \xi$. The physical observation clock still uses χ alone. All three financial channels, their separate coefficients and the prescribed physical-cycle construction remain in use. The physical-driver and magic- ξ variants and their calibrated results are preserved.

Drivers, clocks and observable states. The two independent drivers satisfy

$$\begin{aligned} d\chi_\tau &= (\theta_\chi - \kappa_\chi \chi_\tau) d\tau + \sqrt{2\chi_\tau} dW_\tau^\chi, & \chi_0 &= 0, \\ d\xi_\tau &= (\theta - \kappa \xi_\tau) d\tau + \sqrt{2\xi_\tau} dW_\tau^\xi, & \xi_0 &\geq 0, \\ d\langle \chi, \xi \rangle &= 0. \end{aligned} \tag{4.304}$$

The user-requested public calibration parameters are the sum characteristic q and its jump-driver weight w :

$$q = \theta_\Sigma = \theta_\chi + \theta, \quad w = \frac{\theta_\chi}{q}, \quad \theta_\chi = wq, \quad \theta = (1 - w)q.$$

The entries `theta_sum` and `theta_weight` replace the individual-characteristic descriptor entries; they do not add parameters. Both use change scale 0.05 and bounds $[0, 1]$. The pricing core retains its canonical θ, θ_χ dynamics and converts the public parameters at the pricing boundary. Saved results record the canonical values and the explicit public fit parameters. Then

$$Y = \chi + \xi, \quad \mathbb{E}[dY] = (\theta_\Sigma - \kappa_\chi \chi - \kappa \xi) d\tau, \quad d\langle Y \rangle = 2Y d\tau. \tag{4.305}$$

For unequal reversion rates this drift depends on the composition of Y , so the sum is not a standalone affine process. At $\kappa_\chi = \kappa$ it is affine with characteristic θ_Σ . The covariance identity holds for unequal rates too.

The valid split requires $0 < q \leq 1$, $0 \leq w < 1$ and $0 < (1 - w)q < 1$. At $w = 0$ the physical clock is the jump-free direct identity. Values $q = 0$ or $w = 1$ give zero ξ shape and are invalid; the joint endpoint $q = 1, w = 0$ gives $\theta = 1$ and violates

the retained strict ξ -shape implementation domain. For an active joint boundary clock require $0 < \theta_\Sigma = q < 1$ and define

$$\mathfrak{r}_\tau = \lim_{\epsilon \downarrow 0} \frac{\theta_\Sigma(1 - \theta_\Sigma)}{\epsilon^{\theta_\Sigma}} \int_0^\tau \mathbf{1}_{\{\chi_s + \xi_s < \epsilon\}} ds. \quad (4.306)$$

It increases on the simultaneous boundary $\chi = \xi = 0$. The requested closure $\theta_\Sigma = 1$ has no active joint boundary clock: set its rough increment to zero directly, rather than evaluating a normalized zero divided by zero. Values $\theta_\Sigma > 1$ are excluded. No arbitrary additional margin below one is imposed. When both initial states and reversion rates are zero,

$$\mathbb{E}[\mathfrak{r}_\tau] = \frac{\tau^{1-\theta_\Sigma}}{\Gamma(\theta_\Sigma)}. \quad (4.307)$$

This is not a clock-distribution formula for unequal nonzero rates.

Retain the normalized physical boundary clock of χ from (4.243):

$$\begin{aligned} \mathfrak{r}_\tau^\chi &= \lim_{\epsilon \downarrow 0} \frac{\theta_\chi(1 - \theta_\chi)}{\epsilon^{\theta_\chi}} \int_0^\tau \mathbf{1}_{\{\chi_s < \epsilon\}} ds, \\ C_\tau &= \frac{\Gamma(\theta_\chi)}{(1 - \theta_\chi)\kappa_\chi^{\theta_\chi}} \mathfrak{r}_\tau^\chi, \quad t = C_\tau, \quad \tau(u) = \inf\{s : C_s > u\}. \end{aligned} \quad (4.308)$$

Use $0 < \theta_\chi < 1$, $\kappa_\chi > 0$ for its active boundary clock, and $0 < \theta < 1$, $\kappa \geq 0$ for the retained driver. At $\theta_\chi = 0$, use the existing direct convention $C_\tau = \tau$, $\chi \equiv 0$; the joint clock then reduces to the canonical clock of ξ .

For the configured physical target $u(T)$, the observable states are

$$S_T = S_{\tau(u(T))}, \quad \zeta_T = \zeta_{\tau(u(T))}, \quad \xi_T = \xi_{\tau(u(T))}, \quad \chi_{\tau(u(T))} = 0. \quad (4.309)$$

The ξ state is carried between observations. Only χ is zero at the physical observation boundary. The drivers have independent innovations and retain zero brackets with the financial innovations; the clocks still induce dependence in observed financial states.

Complete financial drift and covariance specification. For the discounted stock, retain the same financial moments as (4.288), with the new clock

meaning:

$$\begin{aligned}
\mathbb{E}_d[dS] &= 0, \\
\frac{d\langle S \rangle}{S^2} &= \frac{d\tau + \boldsymbol{\Psi}_{S,t}dt + \boldsymbol{\Psi}_S d\mathfrak{r}}{\zeta}, \\
\mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\vartheta\boldsymbol{\Psi}_\zeta - \beta\zeta\right) d\tau \\
&\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,t} + \alpha_t - \beta_t\zeta\right) dt \\
&\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,\mathfrak{r}} + \mathfrak{N} - \mathfrak{J}\zeta\right) d\mathfrak{r}, \\
d\langle \zeta \rangle &= \zeta (\boldsymbol{\Psi}_\zeta d\tau + \boldsymbol{\Psi}_{\zeta,t}dt + \boldsymbol{\Psi}_{\zeta,\mathfrak{r}}d\mathfrak{r}), \\
d\langle \zeta, S \rangle &= -S(\rho\sqrt{\boldsymbol{\Psi}_\zeta} d\tau + \rho_t\sqrt{\boldsymbol{\Psi}_{\zeta,t}\boldsymbol{\Psi}_{S,t}} dt \\
&\quad + \mathfrak{N}\sqrt{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}}\boldsymbol{\Psi}_S} d\mathfrak{r}), \\
d\langle \xi, S \rangle &= d\langle \xi, \zeta \rangle = d\langle \xi, \chi \rangle = 0, \\
d\langle \chi, S \rangle &= d\langle \chi, \zeta \rangle = 0.
\end{aligned} \tag{4.310}$$

The magic stock-variance coefficient remains one, fixing the scale of ζ . No Heston or affine-driver stock-variance contribution is added. The signs of ρ , ρ_t and $\mathfrak{N}$ describe variance–stock correlations; reciprocal-state–stock correlations have the opposite sign. The zero driver–stock brackets are specified restrictions of this variant, not consequences of clock singularity.

For positive corresponding reciprocal-state diffusion, the cash and stock-numeraire characteristics remain

$$\begin{aligned}
\vartheta_{d,\tau} &= \vartheta, & \vartheta_{f,\tau} &= \vartheta - \frac{2\rho}{\sqrt{\boldsymbol{\Psi}_\zeta}}, \\
\vartheta_{d,t} &= 1 + \frac{2\alpha_t}{\boldsymbol{\Psi}_{\zeta,t}}, & \vartheta_{f,t} &= 1 + \frac{2(\alpha_t - \rho_t\sqrt{\boldsymbol{\Psi}_{\zeta,t}\boldsymbol{\Psi}_{S,t}})}{\boldsymbol{\Psi}_{\zeta,t}}, \\
\vartheta_{d,\mathfrak{r}} &= 1 + \frac{2\mathfrak{N}}{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}}}, & \vartheta_{f,\mathfrak{r}} &= 1 + \frac{2(\mathfrak{N} - \mathfrak{N}\sqrt{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}}\boldsymbol{\Psi}_S})}{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}}}.
\end{aligned} \tag{4.311}$$

Retain positive ζ_0 , nonnegative diffusion strengths and stock weights, $\vartheta > 1$, the existing physical-channel excess-drift restrictions and positive rough excess for active rough-state diffusion. The retained financial bridge has nonpositive leverage. Apply zero-diffusion drift restrictions directly, without division by a vanishing strength. Both driver characteristics are unchanged under stock numeraire because their stock brackets are zero. The additional joint-clock condition $\theta_\chi + \theta = q \leq 1$ is separate from these reciprocal inverse-moment conditions and is exposed directly by the public sum parameter.

Approved joint launch-clock approximation. Keep the existing prescribed physical grid and observing-driver cycle from (4.261)– (4.266). For each

positive physical increment, its existing conditional physical-clock mean determines a launch duration $\Delta\tau_1$ and deterministic terminal $\chi_1 = \theta_\chi \Delta\tau_1$. This is a prescribed endpoint, not a deterministic sampled launch path: the unresolved affine bridge may revisit zero before that endpoint.

Sample ξ_1 from its retained starting state over the launch duration using the existing native Torch affine transition. Compute the rough launch increment as its conditional mean given both endpoints. Let $p_\chi(s; x, y)$ and $p_\xi(s; x, y)$ be the corresponding affine transition densities and define

$$a_\chi(s) = \frac{1 - e^{-\kappa_\chi s}}{\kappa_\chi}, \quad a_\xi(s) = \frac{1 - e^{-\kappa s}}{\kappa}, \quad (4.312)$$

with $a(s) = s$ at zero reversion. For $0 < \theta_\chi$ and $\theta_\Sigma < 1$, the joint conditional boundary mean is

$$\begin{aligned} & \mathbb{E}[\Delta\tau_1 \mid \chi_0 = 0, \xi_0, \chi_1, \xi_1] \\ &= \frac{1 - \theta_\Sigma}{\Gamma(\theta_\Sigma)} \int_0^{\Delta\tau_1} \frac{\exp[-\xi_0 e^{-\kappa s} / a_\xi(s)]}{a_\chi(s)^{\theta_\chi} a_\xi(s)^\theta} \\ & \quad \times \frac{p_\chi(\Delta\tau_1 - s; 0, \chi_1) p_\xi(\Delta\tau_1 - s; 0, \xi_1)}{p_\chi(\Delta\tau_1; 0, \chi_1) p_\xi(\Delta\tau_1; \xi_0, \xi_1)} ds. \end{aligned} \quad (4.313)$$

The normalization follows from the near-boundary triangular integral

$$\int_{x+y < \epsilon} x^{\theta_\chi - 1} y^{\theta - 1} dx dy = \frac{\Gamma(\theta_\chi) \Gamma(\theta)}{\theta_\Sigma \Gamma(\theta_\Sigma)} \epsilon^{\theta_\Sigma}. \quad (4.314)$$

The two reversion rates remain separate in this kernel. Conditioning on the sampled ξ_1 retains its dependence with the joint clock mean. The calculation still replaces the conditional random clock by a conditional mean; it is not an exact sampler or an exact higher-moment representation.

From the positive launch endpoint, sample the same first-return duration of χ using the existing native Torch Gamma draw. Before that first hit of zero, $\chi > 0$ and hence $Y > 0$. The joint rough clock therefore has exactly zero increment on this final return leg. Advance ξ through that sampled return duration and retain its terminal state. Financial evolution covers the full launch plus return duration. The prescribed conditional-mean physical increment is unchanged. At $\theta_\chi = 0$ use the existing one-driver ξ clock branch directly, and at $\theta_\Sigma = 1$ use zero rough increment.

The extra endpoint transition and joint-density quadrature have been explicitly approved for this variant. They use the established AMD/ROCm Torch backend and do not alter the physical launch proxy, first-return distribution, physical grid or financial step count. The joint-clock kernel, reciprocal core and market adapter are implemented. Their conditional-mean approximation still discards residual clock randomness beyond the sampled endpoints; the small implementation check below is not a quadrature-convergence result.

Combined reciprocal financial step. For each complete slab, combine its three clock increments using

$$\begin{aligned}
A &= \psi_\zeta \Delta\tau + \psi_{\zeta,t} \Delta t + \psi_{\zeta,\mathfrak{r}} \Delta \mathfrak{r}, \\
W &= \Delta\tau + \psi_{S,t} \Delta t + \psi_S \Delta \mathfrak{r}, \\
I &= \frac{1}{2} \vartheta \psi_\zeta \Delta\tau + (\frac{1}{2} \psi_{\zeta,t} + \alpha_t) \Delta t + (\frac{1}{2} \psi_{\zeta,\mathfrak{r}} + \mathfrak{K}) \Delta \mathfrak{r}, \\
R &= \beta \Delta\tau + \beta_t \Delta t + \mathfrak{J} \Delta \mathfrak{r}, \\
B &= \rho \sqrt{\psi_\zeta} \Delta\tau + \rho_t \sqrt{\psi_{\zeta,t} \psi_{S,t}} \Delta t + \mathfrak{M} \sqrt{\psi_{\zeta,\mathfrak{r}} \psi_S} \Delta \mathfrak{r}.
\end{aligned} \tag{4.315}$$

Retain the reference state strength $\psi_\zeta > 0$ and the existing effective parameters

$$D = A/\psi_\zeta, \quad \vartheta_{\text{step}} = 2I/A, \quad \beta_{\text{step}} = R/D, \quad \rho_{\text{step}} = B/\sqrt{AW}. \tag{4.316}$$

The positive magic contribution ensures $A, W > 0$, and Cauchy–Schwarz gives $|\rho_{\text{step}}| \leq 1$. Use the same noncentral-Gamma reciprocal transition and conditional-stock construction, with foreign characteristic $(2I - 2B)/A$. Retain the analytic conditional martingale normalizer, foreign first-state-moment matching and idiosyncratic lognormal residual. The fair inverse-state area is multiplied by W/D to convert from state to stock duration. No sample-dependent adjustment replaces these analytic conditions.

Use exactly the configured number N of financial integration steps per positive physical slab, initially four and finally sixteen. Each receives $\Delta\tau/N$, $\Delta t/N$ and $\Delta \mathfrak{r}/N$, or equivalently state duration D/N and stock duration W/N . No maximum step cap or automatic financial subdivision is added. The unresolved ordering of the three clocks remains the existing frozen-slab approximation. Joint-clock quadrature and financial discretization require separate checks.

Parameter mapping and time dimensions. The public descriptor has twenty model parameters. Its sum and jump-weight entries replace the individual-characteristic entries; the other eighteen keys retain their roles:

Parameter	Descriptor key	Time dimension
ζ_0	<code>zeta_0</code>	$[\tau]$
ϑ	<code>vartheta</code>	dimensionless
β	<code>beta</code>	$[\tau]^{-1}$
Ψ_ζ	<code>evar_zeta</code>	dimensionless
ρ	<code>rho</code>	dimensionless
$q = \theta_\Sigma$	<code>theta_sum</code>	dimensionless
κ_χ	<code>kappa_chi</code>	$[\tau]^{-1}$
$\Psi_{S,t}$	<code>evar_S_physical</code>	$[\tau]/[t]$
$\Psi_{\zeta,t}$	<code>evar_zeta_physical</code>	$[\tau]/[t]$
α_t	<code>alpha_physical</code>	$[\tau]/[t]$
β_t	<code>beta_physical</code>	$[t]^{-1}$
ρ_t	<code>rho_physical</code>	dimensionless
$w = \theta_\chi/q$	<code>theta_weight</code>	dimensionless
κ	<code>kappa</code>	$[\tau]^{-1}$
ξ_0	<code>xi_0</code>	$[\tau]$
Ψ_S	<code>evar_S</code>	$[\tau]^{\theta_\Sigma}$
$\Psi_{\zeta,\mathfrak{r}}$	<code>evar_zeta_rough</code>	$[\tau]^{\theta_\Sigma}$
$\aleph$	<code>aleph</code>	$[\tau]^{\theta_\Sigma}$
$\beth$	<code>bet</code>	$[\tau]^{\theta_\Sigma-1}$
$\mathfrak{M}$	<code>ecorr</code>	dimensionless

The canonical pricing values `theta_chi` = wq and `theta` = $(1 - w)q$ are derived rather than additional editable entries. The optional dated `time_speed` dictionary retains its existing positive interval map, with unit defaults and at least one fixed anchor. Its factors convert calendar increments to model physical increments; they do not rescale the generated magic clock.

The joint clock has units $[\mathfrak{r}] = [\tau]^{1-q}$. Consequently the rough financial strengths depend on the sum characteristic rather than on θ alone. At the zero-state, zero-reversion scaling limit,

$$\mathfrak{r} \sim \tau^{1-q}, \quad t \sim \tau^{1-wq}, \quad \mathfrak{r} \sim t^{(1-q)/(1-wq)}. \quad (4.317)$$

Half this last power is the corresponding clocked-Brownian scaling power. This comparison does not assign a classical local Hurst exponent to the observed stock with jumps. Positive initial state, reversion and a fitted physical map introduce crossover scales.

Implementation, validation and calibration status. The separate `magic_sum_rough_reciproca` kernel, core and market adapter are implemented; seven core and fourteen market tests passed on CPU. A labelled implementation check on the shared Jupyter

server used the AMD Radeon RX 7900 XT, 4,096 paths, four physical intervals, exactly four financial steps and controls disabled. Its raw forward discrepancies were within approximately 1.04 reported Monte Carlo standard errors. Physical-cycle generation took 0.01510 seconds, the joint clock 0.24383 seconds, the previous magic- ξ clock in the timing comparison 0.05029 seconds, and financial integration 0.28443 seconds. These measurements are not a fitted or final result, nor evidence of convergence or full-size pricing performance.

The user paused the previous fit and requested calibration of this variant. Preparation retains the Monday 5 October 2026 EOD snapshot, its nine expirations through 16 October and all 2,517 positive-volume quotes, with the reported close target and existing volume/spread weights. Only the Friday 9 October to Monday 12 October physical interval speed is free; it starts at one, with all other speeds fixed at one. The same staged LM/COBYQA procedure is retained, with the user-requested public sum/weight descriptor above. The separate constrained runtime proposal was discarded before calibration. There is no new numerical margin, hidden projection or adjusted pricing target. Its own checkpoints and notebook remain distinct. No previous variant's score is presented as a result of this model; no new calibration score, variance table or numerical-convergence conclusion is asserted here.

Affine stock contributions with a physical-time rough driver

Extend (4.280) separately, preserving its three reciprocal channels. The ξ affine stock contribution runs in physical time, while the observing χ contribution runs in magic time. Neither the original physical-driver model nor its retained calibration is replaced. The strengths use the $\mathcal{H}/2$ convention of (4.57), with two independent component correlations.

Full cash-measure moments and clocks. For discounted stock, the complete specification is

$$\begin{aligned}
\mathbb{E}_d[dS] &= 0, \\
\frac{d\langle S \rangle}{S^2} &= \frac{d\tau + \mathfrak{V}_{S,t}dt + \mathfrak{V}_S d\mathfrak{T}}{\zeta} + \frac{\mathcal{H}_\xi \xi}{2} dt + \frac{\mathcal{H}_\chi \chi}{2} d\tau, \\
\mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\vartheta \mathfrak{V}_\zeta - \beta \zeta\right) d\tau + \left(\frac{1}{2}\mathfrak{V}_{\zeta,t} + \alpha_t - \beta_t \zeta\right) dt \\
&\quad + \left(\frac{1}{2}\mathfrak{V}_{\zeta,\mathfrak{T}} + \mathfrak{N} - \mathfrak{J} \zeta\right) d\mathfrak{T}, \\
d\langle \zeta \rangle &= \zeta (\mathfrak{V}_\zeta d\tau + \mathfrak{V}_{\zeta,t} dt + \mathfrak{V}_{\zeta,\mathfrak{T}} d\mathfrak{T}), \\
d\langle \zeta, S \rangle &= -S(\rho \sqrt{\mathfrak{V}_\zeta} d\tau + \rho_t \sqrt{\mathfrak{V}_{\zeta,t} \mathfrak{V}_{S,t}} dt \\
&\quad + \mathfrak{D} \sqrt{\mathfrak{V}_{\zeta,\mathfrak{T}} \mathfrak{V}_S} d\mathfrak{T}), \\
\mathbb{E}_d[d\xi] &= (\theta - \kappa \xi) dt, \quad d\langle \xi \rangle = 2\xi dt, \\
\mathbb{E}_d[d\chi] &= (\theta_\chi - \kappa_\chi \chi) d\tau, \quad d\langle \chi \rangle = 2\chi d\tau, \\
d\langle \xi, S \rangle &= \rho_\xi S \sqrt{\mathcal{H}_\xi} \xi dt, \\
d\langle \chi, S \rangle &= \rho_\chi S \sqrt{\mathcal{H}_\chi} \chi d\tau, \\
d\langle \xi, \chi \rangle &= d\langle \xi, \zeta \rangle = d\langle \chi, \zeta \rangle = 0.
\end{aligned} \tag{4.318}$$

The rough clock remains (4.278), the canonical boundary clock of ξ_t in physical time. In magic-time equations, $d\mathfrak{T}$ means $d(\mathfrak{T}_{C_\tau})$. Retain $t = C_\tau = a_\chi \mathfrak{T}_\tau^\chi$ and the cash-normalized physical clock from (4.297). Thus ξ and its rough clock are frozen during each skipped χ excursion, where $dt = 0$. Observations are

$$S_T = S_{\tau(u(T))}, \quad \zeta_T = \zeta_{\tau(u(T))}, \quad \xi_T = \xi_{u(T)}, \quad \chi_{\tau(u(T))} = 0. \tag{4.319}$$

Keep $\chi_0 = 0$ and carry ξ_0 . At $\theta_\chi = 0$, the existing direct branch has $C_\tau = \tau$ and $\chi = 0$.

Each new stock channel has Brownian loading $\sqrt{\mathcal{H}_j j/2}(\rho_j dW^j + \sqrt{1 - \rho_j^2} dW^{j,\perp})$ in its own driver clock: dt for $j = \xi$ and $d\tau$ for $j = \chi$. Therefore the new correlations have their stated component meanings, without another factor 1/2 or $\sqrt{2}$ in their brackets. The independent residuals and reciprocal innovations give a positive semidefinite covariance for independently chosen component correlations in $[-1, 1]$; there is no sum-of-squares restriction.

Numeraire change and parameter units. The affine characteristics remain unchanged, while

$$\begin{aligned}
\mathbb{E}_f[d\xi] &= \left[\theta - (\kappa - \rho_\xi \sqrt{\mathcal{H}_\xi}) \xi\right] dt, \\
\mathbb{E}_f[d\chi] &= \left[\theta_\chi - (\kappa_\chi - \rho_\chi \sqrt{\mathcal{H}_\chi}) \chi\right] d\tau.
\end{aligned} \tag{4.320}$$

The reciprocal stock drift is still (4.300), with this physical-boundary meaning of $d\mathfrak{t}$. Its cash and stock inverse-moment conditions remain (4.281). Keep a_χ at its cash value under the stock measure; replacing its cash κ_χ by the shifted rate would change the observation rule. The stock-measure return-duration law is instead represented by (4.302).

Retain all twenty old parameter meanings and add `H_xi`, `H_chi`, `rho_xi`, `rho_chi`. Their units distinguish the two driver clocks:

$$\begin{aligned} [\xi] &= [t], & [\zeta] &= [\chi] = [\tau], \\ [\mathcal{H}_\xi] &= [t]^{-2}, & [\mathcal{H}_\chi] &= [\tau]^{-2}, \\ [\mathfrak{t}] &= [t]^{1-\theta}, & [\mathfrak{v}_S] &= [\mathfrak{v}_{\zeta, \mathfrak{t}}] = [\mathfrak{N}] = [\tau]/[t]^{1-\theta}, \\ [\kappa] &= [t]^{-1}, & [\beth] &= [t]^{\theta-1}. \end{aligned} \tag{4.321}$$

Both correlations are dimensionless. In common engineering time units these reduce to the previous power rules; their boundary scaling must still distinguish physical time from magic time.

Separate finite-grid construction and status. Keep the original prescribed physical cycles, physical ξ rough bank and combined reciprocal financial slabs. For ξ , reuse the retained outer physical endpoints and normalized Gamma endpoint factor with the foreign affine mean. Evaluate (4.303) at $L = \Delta t$, not $\Delta\tau$; its independent residual is $(1 - \rho_\xi^2)\mathcal{H}_\xi/2$ times this area.

For χ , reuse the first-return construction described after (4.302): retain exactly the configured N complete-cycle financial subdivisions, positive conditional return areas, conditional endpoint factors and the cash-to-stock duration likelihood. The prescribed deterministic launch receives its independent compensated Gaussian variance and discards launch leverage. Observed zero endpoints do not remove its return area or duration likelihood. These analytic discrete normalizers do not make the unresolved joint driver, rough-clock and stock path law exact. Mean-area proxies must not be identified with observed jump quadratic variation without further analysis.

Setting both added strengths to zero delegates directly to the original physical-driver simulator, preserving its random inputs and outputs. The existing reciprocal conditions and helper domains remain: the initial implementation supports negative component leverage, nonnegative cash affine rates and nonnegative shifted χ return rate, and retains the endpoint tilt's finite-second-moment check. These helper restrictions are distinct from continuous covariance admissibility and add no arbitrary numerical margin.

The separate model identifier is `magic_physical_heston_reciprocal_v1`; its market adapter adds `_prescribed_physical_cycles_v1`. Its separate core, mar-

ket and calibration adapters and notebook are implemented, with twenty-one focused CPU tests passing. The three-extension comparison uses the same nine October expirations, one Friday-to-Monday speed and existing joint-trial, residual-secant and SVD search; it does not replace that search by a generic Python optimizer. The separate notebooks record current calibration scores, fitted parameters and raw million-path/N16 confirmations. Provisional controlled fitting remains distinct. Passing implementation tests or completing a fitting pass does not establish convergence or an optimal fit.

Affine stock contributions with a joint magic-time boundary clock

Extend (4.310) separately with the same two affine stock channels as (4.295). Both drivers and both added stock contributions run in magic time. The original sum-clock model and its retained calibration remain intact. The physical boundary clock uses χ alone, while the rough clock uses the joint boundary of $Y = \chi + \xi$.

Full moments and the joint clock. The complete cash specification is

$$\begin{aligned}
\mathbb{E}_d[dS] &= 0, \\
\frac{d\langle S \rangle}{S^2} &= \frac{d\tau + \boldsymbol{\Psi}_{S,t}dt + \boldsymbol{\Psi}_S d\mathfrak{r}}{\zeta} + \frac{\mathcal{H}_\xi \xi + \mathcal{H}_\chi \chi}{2} d\tau, \\
\mathbb{E}_d[d\zeta] &= \left(\frac{1}{2}\vartheta \boldsymbol{\Psi}_\zeta - \beta \zeta\right) d\tau + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,t} + \alpha_t - \beta_t \zeta\right) dt \\
&\quad + \left(\frac{1}{2}\boldsymbol{\Psi}_{\zeta,\mathfrak{r}} + \mathfrak{N} - \mathfrak{J}\zeta\right) d\mathfrak{r}, \\
d\langle \zeta \rangle &= \zeta (\boldsymbol{\Psi}_\zeta d\tau + \boldsymbol{\Psi}_{\zeta,t}dt + \boldsymbol{\Psi}_{\zeta,\mathfrak{r}}d\mathfrak{r}), \\
d\langle \zeta, S \rangle &= -S(\rho\sqrt{\boldsymbol{\Psi}_\zeta}d\tau + \rho_t\sqrt{\boldsymbol{\Psi}_{\zeta,t}\boldsymbol{\Psi}_{S,t}}dt \\
&\quad + \mathfrak{N}\sqrt{\boldsymbol{\Psi}_{\zeta,\mathfrak{r}}\boldsymbol{\Psi}_S}d\mathfrak{r}), \\
\mathbb{E}_d[d\xi] &= (\theta - \kappa\xi)d\tau, \quad d\langle \xi \rangle = 2\xi d\tau, \\
\mathbb{E}_d[d\chi] &= (\theta_\chi - \kappa_\chi\chi)d\tau, \quad d\langle \chi \rangle = 2\chi d\tau, \\
d\langle \xi, S \rangle &= \rho_\xi S \sqrt{\mathcal{H}_\xi} \xi d\tau, \\
d\langle \chi, S \rangle &= \rho_\chi S \sqrt{\mathcal{H}_\chi} \chi d\tau, \\
d\langle \xi, \chi \rangle &= d\langle \xi, \zeta \rangle = d\langle \chi, \zeta \rangle = 0.
\end{aligned} \tag{4.322}$$

The new correlations have the component meaning in (4.296). Independent residual innovations give a positive semidefinite full covariance without a sum-of-squared-correlations constraint. All reciprocal coefficients, correlations and drift conventions are preserved.

Keep the public characteristic and jump weight

$$q = \theta_\chi + \theta, \quad w = \theta_\chi/q, \quad \theta_\chi = wq, \quad \theta = (1 - w)q. \quad (4.323)$$

The editable keys `theta_sum`, `theta_weight` replace the two individual shapes, with the established scales and bounds. Require $0 < q \leq 1$, $0 \leq w < 1$, $0 < (1 - w)q < 1$. Active rough time uses $q < 1$ and the canonical clock

$$\mathfrak{r}_\tau = \lim_{\epsilon \downarrow 0} \frac{q(1 - q)}{\epsilon^q} \int_0^\tau \mathbf{1}_{\{\chi_s + \xi_s < \epsilon\}} ds. \quad (4.324)$$

At $q = 1$ set its rough increment to zero directly. At $w = 0$ retain the exact physical-clock identity; $q = 0$ or $w = 1$ would give an invalid zero ξ shape. Retain the cash physical clock (4.308) and observed states (4.309), with $\chi_0 = 0$ and retained ξ_0 . No additional state or clock is introduced.

Stock measure and the sum characteristic. The separate affine rate shifts are (4.299); the reciprocal stock drift is (4.300) with the joint-clock meaning of $d\mathfrak{r}$. The sum has

$$\begin{aligned} d\langle Y \rangle &= 2Y d\tau, \\ \mathbb{E}_f[dY] &= (q - \kappa_\chi^\mathfrak{f} \chi - \kappa^\mathfrak{f} \xi) d\tau, \\ \frac{d\langle Y, S \rangle}{S} &= (\rho_\xi \sqrt{\mathcal{H}_\xi} \xi + \rho_\chi \sqrt{\mathcal{H}_\chi} \chi) d\tau. \end{aligned} \quad (4.325)$$

Both individual characteristics, and therefore q, w , are invariant. For unequal rates the sum is not a standalone affine process. Equal cash rates need not remain equal after independent stock-measure shifts. The local-time normalization uses the same q under both numeraires. Likewise, keep the physical-clock multiplier fixed at its cash value; the stock return-duration law changes through its density likelihood, not through renormalizing physical time.

Keep the old public twenty entries and add `H_xi`, `H_chi`, `rho_xi`, `rho_chi`. The intrinsic units are

$$\begin{aligned} [\xi] &= [\chi] = [\zeta] = [\tau], & [\mathcal{H}_\xi] &= [\mathcal{H}_\chi] = [\tau]^{-2}, \\ [\mathfrak{r}] &= [\tau]^{1-q}, & [\Psi_S] &= [\Psi_{\zeta, \mathfrak{r}}] = [\mathfrak{N}] = [\tau]^q, \\ [\mathfrak{z}] &= [\tau]^{q-1}. \end{aligned} \quad (4.326)$$

The shapes, weight and correlations are dimensionless. Relative to the physical boundary scaling, the rough-clock power remains $(1 - q)/(1 - \theta_\chi)$; the affine additions do not change this intrinsic clock comparison or identify a local Hurst exponent for the observed process with jumps.

Retained approximation, nesting and status. Reuse the approved joint launch-clock conditional mean in (4.313), conditioned on the same sampled ξ launch endpoint. Sample its ξ return endpoint with the existing affine transition over the retained χ return duration. The joint rough increment stays zero during that first return: $\chi > 0$ before its final hit. The conditional clock mean still discards remaining clock randomness and is not an exact joint-clock sampler.

The composed launch and return transitions give the outer ξ CIR marginal. Apply the existing normalized endpoint stock factor with its shifted affine mean, and the positive forward-mean area in (4.303) at $L = \Delta\tau$. The residual variance retains $(1 - \rho_\xi^2)\mathcal{H}_\xi/2$ times that area. This does not reproduce every joint rough-clock, area and stock fluctuation.

Reuse the same χ conditional return transitions, positive areas and duration likelihood from (4.302), with exactly N financial subdivisions of every complete cycle. The prescribed deterministic launch has independent compensated Gaussian stock variance and discards launch leverage. All old reciprocal slab coefficients and analytic normalizers remain unchanged; no sample forward correction or clock rescaling is added.

At $\mathcal{H}_\xi = \mathcal{H}_\chi = 0$, delegate to the old sum simulator bitwise, preserving draws and snapshots. At $\theta_\chi = 0$, the χ contribution vanishes. Keep the old reciprocal inverse-moment conditions and the stated native helper domains, including non-negative return rates, negative implementation leverage and finite endpoint-tilt second moments; these do not add an arbitrary new financial bound.

The separate model identifier is `magic_sum_heston_reciprocal_v1`; its market adapter adds `_prescribed_physical_cycles_v1`. Its separate core, market and calibration adapters and notebook are implemented, with twenty CPU tests passing. Rotation with the two other affine extensions uses the same nine-expiry target, single weekend speed and authorized coupled joint-trial, residual-secant and SVD search. Raw million-path/N16 confirmation remains distinct from provisional controlled fitting. Current errors, parameter vectors and stage counts are recorded in the separate notebooks and saved results. Passing implementation tests or completing a fitting pass does not establish convergence or an optimal fit.

An affine-sum volatility model in magic time

Consider a separate specification with two affine states, ξ and Ξ , in magic time τ . The process Ξ generates the physical boundary clock and also enters stock variance. There is no separate reciprocal state ζ and no additional rough clock. The reciprocal variance contribution is the inverse of $\xi + \Xi$; two linear variance contributions permit separate stock correlations with the two drivers. The symbols K and P are upper-case parameter commands rendered as the Roman glyphs K

and P , distinct from κ and ρ .

Cash-numeraire moments. For the discounted stock, the implemented affine baseline moments are

$$\begin{aligned}
\mathbb{E}_d[dS_\tau] &= 0, \\
\mathbb{E}_d[d\xi_\tau] &= (\theta - \kappa\xi_\tau)d\tau, \\
\mathbb{E}_d[d\Xi_\tau] &= (\Theta - K\Xi_\tau)d\tau, \\
d\langle\xi\rangle_\tau &= 2\xi_\tau d\tau, & d\langle\Xi\rangle_\tau &= 2\Xi_\tau d\tau, \\
d\langle\xi, \Xi\rangle_\tau &= 0, \\
\frac{d\langle S\rangle_\tau}{S_\tau^2} &= \left[\frac{\aleph}{\xi_\tau + \Xi_\tau} + \Sigma\Xi_\tau + \sigma\xi_\tau \right] d\tau, \\
\frac{d\langle\xi, S\rangle_\tau}{S_\tau} &= \rho\xi_\tau\sqrt{2\sigma} d\tau, \\
\frac{d\langle\Xi, S\rangle_\tau}{S_\tau} &= P\Xi_\tau\sqrt{2\Sigma} d\tau.
\end{aligned} \tag{4.327}$$

The nonnegative strengths are $\aleph, \Sigma, \sigma$ and the two correlation parameters satisfy $|\rho|, |P| \leq 1$. A Brownian realization uses independent $W^\xi, W^\Xi, W^\perp$ and

$$\begin{aligned}
\frac{dS_\tau}{S_\tau} &= \rho\sqrt{\sigma\xi_\tau} dW_\tau^\xi + P\sqrt{\Sigma\Xi_\tau} dW_\tau^\Xi \\
&\quad + \sqrt{\frac{\aleph}{\xi_\tau + \Xi_\tau} + (1 - \rho^2)\sigma\xi_\tau + (1 - P^2)\Sigma\Xi_\tau} dW_\tau^\perp, \\
d\xi_\tau &= (\theta - \kappa\xi_\tau)d\tau + \sqrt{2\xi_\tau} dW_\tau^\xi, \\
d\Xi_\tau &= (\Theta - K\Xi_\tau)d\tau + \sqrt{2\Xi_\tau} dW_\tau^\Xi.
\end{aligned} \tag{4.328}$$

Thus $\Sigma\Xi$ and $\sigma\xi$ are Heston-type linear variance channels. The independent stock variance left after the correlated channels is exactly the nonnegative expression under the final square root. The parameters ρ and P are correlations with their respective linear stock channels; the correlation with the complete stock innovation is attenuated by the other variance contributions.

The retained parameter domain is

$$0 < \theta < 1, \quad 0 < \Theta < 1, \quad \theta + \Theta > 1, \quad \xi_0 > 0, \quad \Xi_0 = 0. \tag{4.329}$$

Each individual affine state may reach zero, while their sum has a Feller characteristic greater than one. Near the joint origin, the inverse-sum moment at a positive deterministic magic time has the integrability condition $\theta + \Theta > 1$. This does not give an integrability bound for $1/\xi$ at physical boundary observations.

Reuse of the reciprocal model. The sum of the two independent affine states has moments

$$\begin{aligned}\mathbb{E}[d(\xi + \Xi)] &= (\theta + \Theta - \kappa\xi - K\Xi)d\tau, \\ d\langle \xi + \Xi \rangle &= 2(\xi + \Xi)d\tau.\end{aligned}\tag{4.330}$$

If $\kappa = K$, its unconditional scalar dynamics are precisely the reciprocal model's denominator process, with state $\xi + \Xi$, characteristic $\theta + \Theta$, diffusion strength 2, and the common reversion coefficient. The existing fair-integrated-variance formula can therefore be reused for its unconditional fixed-magic-time law, multiplied by $\aleph$. Unequal reversions retain dependence on the state split; the scalar sum is then not an affine process with constant reversion. No equality of these two parameters is imposed here.

Conditioning Ξ on a sampled first return, or prescribing its launch, also changes the aggregate transition law. Even at equal reversions, the unconditional reciprocal formula is not thereby an exact conditional return-bridge variance formula. The existing analytic normalization of stock endpoint factors can nevertheless be reused on each actual conditional Gamma transition. The reciprocal variance channel itself has no state-stock covariance in (4.327); leverage comes from the two linear variance channels.

Stock numeraire. The covariance shifts only the proportional drifts. Under the stock numeraire the affine characteristics remain θ and Θ , and the two reversion coefficients become

$$\kappa^S = \kappa - \rho\sqrt{2\sigma}, \quad K^S = K - P\sqrt{2\Sigma}.\tag{4.331}$$

These shifted coefficients are properties of the model's change of numeraire, rather than corrections based on simulated stock means. The reciprocal stock channel contributes no state-stock covariance. Because Ξ shares financial noise, its return-duration law under stock weighting differs from its cash-numeraire law when $P\sqrt{\Sigma} \neq 0$. It cannot be treated as an independent observing clock in the stock conditional moments.

Rational reciprocal leverage and a common numeraire family. The agreed complementary extension retains positive constant immigration for both states and allocates the reciprocal stock channel along their sum. Introduce a

common rational drift coefficient α and reciprocal leverage coefficient $\mathfrak{N}$:

$$\begin{aligned}
\mathbb{E}_d[d\xi] &= \left[\theta + \alpha \frac{\xi}{\xi + \Xi} - \kappa \xi \right] d\tau, \\
\mathbb{E}_d[d\Xi] &= \left[\Theta + \alpha \frac{\Xi}{\xi + \Xi} - K\Xi \right] d\tau, \\
\frac{d\langle \xi, S \rangle}{S} &= \left[\rho \xi \sqrt{2\sigma} + \mathfrak{N} \sqrt{2\aleph} \frac{\xi}{\xi + \Xi} \right] d\tau, \\
\frac{d\langle \Xi, S \rangle}{S} &= \left[P\Xi \sqrt{2\Sigma} + \mathfrak{N} \sqrt{2\aleph} \frac{\Xi}{\xi + \Xi} \right] d\tau.
\end{aligned} \tag{4.332}$$

The stock and state quadratic variations and $d\langle \xi, \Xi \rangle = 0$ in (4.327) are retained. The two rational fractions add to one. The additional covariance with the sum is therefore $\mathfrak{N} \sqrt{2\aleph} d\tau$, and $\mathfrak{N}$ is the reciprocal channel's correlation with the radial Brownian motion of $\xi + \Xi$. Its individual state correlations are $\mathfrak{N} \sqrt{\xi/(\xi + \Xi)}$ and $\mathfrak{N} \sqrt{\Xi/(\xi + \Xi)}$.

The cash and stock drifts have exactly the same parametric form:

$$\begin{aligned}
\mathbb{E}^S[d\xi] &= \left[\theta + \alpha^S \frac{\xi}{\xi + \Xi} - \kappa^S \xi \right] d\tau, \\
\mathbb{E}^S[d\Xi] &= \left[\Theta + \alpha^S \frac{\Xi}{\xi + \Xi} - K^S \Xi \right] d\tau, \\
\alpha^S &= \alpha + \mathfrak{N} \sqrt{2\aleph}, & \kappa^S &= \kappa - \rho \sqrt{2\sigma}.
\end{aligned} \tag{4.333}$$

The constants θ, Θ remain unchanged; K^S retains the shift in (4.331). Multiplying the entire immigration θ by the ratio would remove the positive drift at $\xi = 0, \Xi > 0$, allowing absorption and invalidating the previous joint-origin argument. The added term instead vanishes there.

The aggregate moments are now

$$\begin{aligned}
\mathbb{E}_d[d(\xi + \Xi)] &= [\theta + \Theta + \alpha - \kappa \xi - K\Xi] d\tau, \\
\mathbb{E}^S[d(\xi + \Xi)] &= [\theta + \Theta + \alpha^S - \kappa^S \xi - K^S \Xi] d\tau, \\
d\langle \xi + \Xi \rangle &= 2(\xi + \Xi) d\tau.
\end{aligned} \tag{4.334}$$

The positive-magic-time inverse-sum conditions in the two numeraires are

$$\theta + \Theta + \alpha > 1, \quad \theta + \Theta + \alpha^S > 1. \tag{4.335}$$

These are conditions protecting the joint origin; they do not by themselves prove the stock martingale property after physical-time observation. Exact scalar affine

closure requires $\kappa = K$ in cash measure and, separately, $\kappa^S = K^S$ in stock measure. Equal cash reversions do not imply equal stock reversions.

The independent stock-variance rate must account for the reciprocal contribution sharing both state Brownian motions. It is

$$(1 - \mathfrak{N}^2) \frac{\aleph}{\xi + \Xi} + (1 - \rho^2) \sigma \xi + (1 - P^2) \Sigma \Xi - 2\mathfrak{N}\sqrt{\aleph} \frac{\rho\sqrt{\sigma}\xi + P\sqrt{\Sigma}\Xi}{\xi + \Xi}. \quad (4.336)$$

Its nonnegativity is a covariance requirement. In particular, $\rho, P \leq 0$, $\mathfrak{N} \geq 0$ and $|\rho|, |\mathfrak{N}|, |P| \leq 1$ make every displayed term nonnegative. Positive $\mathfrak{N}$ gives negative stock-reciprocal-volatility leverage, because increasing ξ decreases the reciprocal variance channel; the negative ρ gives negative leverage for the linear $\sigma\xi$ channel.

The baseline implementation described below corresponds to $\alpha = \mathfrak{N} = 0$. Nonzero α makes the individual cash transitions coupled and nonaffine. In particular, the physical Ξ return is no longer the unchanged affine Gamma law. Even at $\alpha = 0$, nonzero $\mathfrak{N}$ changes its stock-weighted first-return law, so the baseline duration-density ratio cannot be reused unchanged. We consider the integrating-factor construction below. Its numerical physical-clock matching and validated stock compensation remain pending; the previously reported calibration scores belong to the baseline.

First calibration stage: affine cash dynamics with reciprocal leverage. The separate initial calibration fixes $\alpha = 0$ and allows $\mathfrak{N}$ to vary. This is a subfamily of (4.332), not a fit of the coupled nonlinear cash clock. Its stock-numeraire rational drift remains $\alpha^S = \mathfrak{N}\sqrt{2\aleph}$ and need not vanish. Cash states, the prescribed physical launch and return, the configured full-cycle step count, and the analytically normalized linear stock factors are retained. The linear first-return density factor is also retained for its linear contribution alone.

Only the rational stock contribution is added, using the potential $\mathfrak{N}\sqrt{\aleph/2} \log(\xi + \Xi)$, its cash Itô drift and the linear-rational covariance compensation. Each step reuses the existing reciprocal inverse-area estimate and freezes its beginning-state fractions. The rational squared exposure and twice its cross exposure are subtracted from the previous independent stock-variance area. The deterministic launch has no Ξ Brownian exposure, while its independent $\Sigma\Xi$ stock noise is retained. These are finite-step Euler approximations; they do not establish exact finite-step stock normalization. The rational return weighting is included in this path factor, without adding another return-density ratio or any sample normalizer. Raw stock forwards and put-call parity errors must be reported at every expiry. At $\mathfrak{N} = 0$ the separate implementation delegates to the baseline exactly.

The separate notebook is `example_magic_affine_sum_rational_calibration.ipynb`. Its selected SPX/SPXW target contains 33 strips on 30 dates, 6,487 positive-volume quotes, from seven days through three months; both rights and all supplied in- and out-of-the-money strikes retain the established volume/bid/ask objective. The first date is fitted before assessing candidates on the full target. Initial fitting uses 8,192 paths, four magic integration steps per positive one-day physical interval and explicitly labelled control variates. Final pricing requires 2^{20} paths, sixteen configured steps and raw uncontrolled prices. Date-speed nesting, checkpoint resume, live counts and plots are retained. No calibration result for this stage is asserted here yet.

Integrating factor and an intrinsic affine return. Within a physical cycle define an auxiliary positive integrating factor A , initially one, and intrinsic time $\tilde{\tau}$:

$$\begin{aligned} A_\tau &= \exp \left[\int_0^\tau \left(K - \frac{\alpha}{\xi_s + \Xi_s} \right) ds \right], & d\tilde{\tau} &= A_\tau d\tau, \\ U_\xi &= A\xi, & U_\Xi &= A\Xi. \end{aligned} \quad (4.337)$$

Resetting $A = 1$ at an observation changes auxiliary coordinates only; the observed financial states are carried forward. Product Itô calculus gives

$$\begin{aligned} dU_\Xi &= \Theta d\tilde{\tau} + \sqrt{2U_\Xi} dB^\Xi, \\ dU_\xi &= \left[\theta + (K - \kappa) \frac{U_\xi}{A} \right] d\tilde{\tau} + \sqrt{2U_\xi} dB^\xi, \\ \frac{dA}{d\tilde{\tau}} &= K - \frac{\alpha A}{U_\xi + U_\Xi}, & \frac{d\tau}{d\tilde{\tau}} &= \frac{1}{A}. \end{aligned} \quad (4.338)$$

The transformed Brownian motions have zero cross-covariance. The physical driver U_Ξ is an autonomous zero-reversion affine process for unequal as well as equal original reversions. When $\kappa = K$, both transformed states are independent affine processes. From a positive transformed launch endpoint $\tilde{y}$, the intrinsic first-return duration is exactly $\tilde{y}/G$, with $G \sim \text{Gamma}(1 - \Theta, 1)$. Its conditional return bridge retains characteristic $2 - \Theta$. This exact intrinsic law does not give the actual magic return duration without integrating $1/A$.

The stock moments become

$$\begin{aligned}\frac{d\langle S \rangle}{S^2} &= \left[\frac{\aleph}{U_\xi + U_\Xi} + \frac{\sigma U_\xi + \Sigma U_\Xi}{A^2} \right] d\tilde{\tau}, \\ \frac{d\langle U_\xi, S \rangle}{S} &= \left[\rho \frac{U_\xi \sqrt{2\sigma}}{A} + \mathfrak{N} \sqrt{2\aleph} \frac{U_\xi}{U_\xi + U_\Xi} \right] d\tilde{\tau}, \\ \frac{d\langle U_\Xi, S \rangle}{S} &= \left[\mathfrak{P} \frac{U_\Xi \sqrt{2\Sigma}}{A} + \mathfrak{N} \sqrt{2\aleph} \frac{U_\Xi}{U_\xi + U_\Xi} \right] d\tilde{\tau}.\end{aligned}\tag{4.339}$$

The reciprocal variance strength remains $\aleph$; the linear strengths become σ/A^2 and Σ/A^2 .

The canonical boundary clock transforms as

$$d\mathfrak{X}^\Xi = A^{\Theta-1} d\mathfrak{X}^{U_\Xi}.\tag{4.340}$$

Indeed, $\Xi = U_\Xi/A$ and $d\tau = d\tilde{\tau}/A$ turn the occupation limit $\varepsilon^{-\Theta} \int \mathbf{1}_{\{\Xi < \varepsilon\}} d\tau$ into the weighted occupation limit for U_Ξ . The same factor applies with the paper's fixed canonical multiplier. Observations remain $U_\Xi = 0$, but physical clock increments depend on the path of A .

For an intrinsic launch of duration $\tilde{h}$ and endpoint $\tilde{y}$, the zero-rate conditional edge kernel is

$$\begin{aligned}\mathbb{E}\left[d\mathfrak{X}_s^{U_\Xi} \mid U_{\Xi,0} = 0, U_{\Xi,\tilde{h}} = \tilde{y}\right] \\ = \frac{1-\Theta}{\Gamma(\Theta)} s^{-\Theta} \left(\frac{\tilde{h}}{\tilde{h}-s} \right)^\Theta \exp\left[-\tilde{y} \left(\frac{1}{\tilde{h}-s} - \frac{1}{\tilde{h}} \right)\right] ds.\end{aligned}\tag{4.341}$$

For nonzero α , the required physical conditional mean contains $A_s^{\Theta-1}$ inside the expectation. A sampled or prescribed A depends on the same state path; multiplying this kernel by a frozen weight is an approximation. The numerical trial retains the daily physical grid and configured counts 4 and 16. Its launch predictor, finite-step integration and weighted-edge mismatch remain to be specified and tested; prescribed grid labels do not establish exact physical-clock matching.

At $\alpha = 0$, $A = 1 + K\tilde{\tau}$ is deterministic. The existing finite-reversion launch h , with physical-state endpoint $y = \Theta h$, has exact transformed coordinates

$$\tilde{h} = \frac{e^{Kh} - 1}{K}, \quad \tilde{y} = e^{Kh} \Theta h.\tag{4.342}$$

The weighted zero-rate kernel reproduces the existing normalized conditional clock mean. Using $\tilde{y} = \Theta \tilde{h}$ or only the small-step power law would change the finite-reversion baseline. Validate the trial at $\alpha = \mathfrak{N} = 0$ before extending calibration.

In this intrinsic formulation $K = 0$ is mathematically allowed, but the existing normalized physical-clock multiplier $K^{-\Theta}$ is then undefined and cannot be silently retained.

Implemented baseline: physical observations and prescribed grid.

Use the same normalization of the affine boundary clock as in (4.243), now generated by Ξ :

$$C_\tau = \frac{\Gamma(\Theta)}{(1 - \Theta)K^\Theta} \mathfrak{r}_\tau^\Xi, \quad \tau(u) = \inf\{s : C_s > u\}, \quad K > 0. \quad (4.343)$$

The observed states are $S_T = S_{\tau(u(T))}$ and $\xi_T = \xi_{\tau(u(T))}$, with $\Xi = 0$ at physical observations. Start the physical driver at zero on every observed valuation date. The deterministic business-time and fitted interval-speed map $u(T)$ is applied before the inverse clock. Forward curves, discounts and expiration dates retain their market conventions. Expiries sharing a date share the same speed, and at least one speed remains fixed to anchor the deterministic time scale.

For the selected simulation approximation, reuse the established regular mapped physical grid with exact expiration nodes. For a positive physical interval Δu , choose the launch duration h from (4.261) and (4.263), replacing $(\theta_\chi, \kappa_\chi)$ by (Θ, K) . The approved deterministic launch is explicitly

$$\Xi_u^{\text{launch}} = \Theta u, \quad 0 \leq u \leq h, \quad y = \Theta h. \quad (4.344)$$

This is a linear proxy, including at finite mean reversion; it is not the stochastic affine launch path or its exact mean-reverting mean. From y sample the future first return to zero using the existing Gamma law

$$D = \frac{1}{K} \log \left(1 + \frac{Ky}{G} \right), \quad G \sim \text{Gamma}(1 - \Theta, 1). \quad (4.345)$$

The state simulation evolves ξ over the launch and the entire sampled return. Conditional on this first-return duration, the positive Ξ return bridge uses the existing affine dual characteristic $2 - \Theta$, rather than an unconditioned affine endpoint transition. At the next physical node retain ξ , set $\Xi = 0$, and advance physical time by the prescribed Δu . The separate stock implementation evolves its conditional moments over these same intervals.

The fixed integration count remains configurable and applies to every path across its complete launch-plus-return duration $h + D$ per positive physical interval: initially 4 and finally 16. The launch does not add a further financial integration step. The complete launch and return must be covered without dropping long returns or adding a maximum-step refinement. The physical-driver surrogate retains

the launch and conditional-clock approximations already identified in section 4.1; use of its Gamma first-return law does not make the complete clock surrogate exact.

Inverse-sum variance integration. Full variance integration cannot use the instantaneous boundary value $\aleph/\xi$ as a trapezoid endpoint. At physical observation nodes $\Xi = 0$; affine ξ laws with $\theta < 1$ can have $\mathbb{E}[1/\xi] = \infty$, even though the joint positive-magic-time inverse-sum moment is finite. A sufficiently small realized ξ would therefore create a spurious, unbounded endpoint contribution. The reciprocal contribution is instead integrated using the approved local affine approximation described below. The sampled Ξ states retain their first-return conditioning, while the local scalar fair-variance approximation does not reproduce that full bridge law. The reciprocal area must not be replaced by a boundary trapezoid, an arbitrary denominator floor or clipped paths. The two linear channels have no inverse-state boundary singularity.

Completed state simulation. The separate module `py/magic_affine_sum.py` implements parameter validation, the deterministic launch, and the retained ordinary affine and first-return-conditioned affine samplers. The module `py/magic_affine_sum_paths.py` implements complete-cycle state simulation with exactly the configured number of equal subdivisions of $h + D$. A subdivision straddling launch and return remains one step. Its snapshots contain ξ, Ξ, τ ; every physical observation has $\Xi = 0$. Its metadata explicitly reports that stock pricing is not implemented. The ordinary affine sampler supports $\kappa \geq 0$; this is an implementation restriction, not a further mathematical stability condition or an arbitrary finite parameter bound. No calibration or stock-martingale verification is reported for this state-only implementation.

Stock return weighting and the local affine Euler step. Write $f_K(D | y)$ for the exact first-return density associated with (4.345). At the prescribed launch endpoint $y = \Theta h$, stock weighting of a return cycle must include the likelihood ratio

$$w(D | y) = \frac{f_{K^S}(D | y)}{f_K(D | y)}, \quad \mathbb{E}_d[w(D | y)] = 1 \quad \text{when } K^S \geq 0. \quad (4.346)$$

Both densities then integrate to one. Normalizing the stock only inside a domestic return-conditioned bridge would remove this necessary duration tilt. The finite-step construction reuses the existing normalized exponential tilts of Gamma endpoint laws for the ordinary ξ transition and the conditional Ξ return bridge, targeting their stock-numeraire endpoint means. Their zeroth and targeted first moments

are analytic, but mean matching is not the exact finite-step stock-numeraire bridge distribution. For each actually sampled conditional endpoint law, the retained normalization is

$$M_x = \frac{\exp(a_x x_1)}{\mathbb{E}_d[\exp(a_x x_1) \mid \mathcal{I}]}, \quad \mathbb{E}_d[M_x \mid \mathcal{I}] = 1, \quad x \in \{\xi, \Xi\}, \quad (4.347)$$

where $\mathcal{I}$ contains the current states and the retained clock conditioning. The denominator is the existing analytic Gamma Laplace normalizer, not a simulated average. Multiplication by an independent conditional Gaussian factor $\exp(\sqrt{v}Z - v/2)$ preserves this first moment for any finite nonnegative conditional variance v . At a physical-cycle observation, the factor $w(D \mid y)$ must remain outside the product of normalized bridge factors. Its expectation one, combined with the conditional endpoint normalizations, is the required martingale construction. Variance-approximation accuracy is a separate question from this analytic first-moment identity. The separate helper module `py/magic_affine_sum_moments.py` now reuses the existing Gamma normalizer and implements the return-time density ratio. It includes the deterministic final-zero endpoint, whose factor is exactly one. These helpers have analytic moment tests and a raw native-GPU first-moment check. The complete stock module `py/magic_affine_sum_stock.py` now combines these factors with the approved local reciprocal Euler area and linear-channel residual variances. Its conditional snapshots retain the raw log forward, independent log-normal variance, both affine states, magic time, the accumulated fair-area proxy, and the clock likelihood. The separate `py/magic_affine_sum_market.py` adapter uses the existing conditional-lognormal option payoffs and source volume/bid/ask target. It retains every supplied option right and moneyness, raw forward diagnostics, and the prescribed physical grid. No calibrated option-price result is claimed here.

The approved deterministic launch has no random Ξ innovation with which to realize the specified P leverage. Its stock convention supplies the full $\Sigma\Xi$ stock variance as an independent compensated lognormal contribution, with launch variance $\Sigma\Theta h^2/2$ and no launch leverage. This is a departure from the stochastic launch covariance, not an identity of the continuous model.

The approved reciprocal area approximation treats $\xi + \Xi$ as locally affine, freezing the split-dependent reversion at the start of each configured Euler step:

$$\kappa_{\text{eff}} = \frac{\kappa\xi + K\Xi}{\xi + \Xi}. \quad (4.348)$$

Reuse the existing reciprocal fair-integrated-variance formula with initial state $\xi + \Xi$, Feller characteristic $\theta + \Theta$, diffusion strength 2, and rate κ_{eff} , and multiply its result by the stock strength $\aleph$. If $\mathcal{V}$ denotes that existing expected integrated

inverse-state formula, the selected Euler-step contribution is

$$\Delta v_{\text{reciprocal}} = \aleph \mathcal{V}(\Delta\tau; \xi + \Xi, \theta + \Theta, \kappa_{\text{eff}}, 2). \quad (4.349)$$

This integrates the frozen affine reciprocal mean within the step, instead of multiplying a potentially singular boundary value by its duration. The formula is exact for the unconditional aggregate at equal reversion rates; freezing unequal rates, imposing a deterministic launch, and conditioning the physical driver on its future first return make its application here a local Euler approximation. The configured N steps still cover each complete $h + D$ cycle. Analytic stock-factor normalization preserves the conditional mean in this construction; variance accuracy and market calibration require separate validation.

Stored variance and the clock likelihood. The implementation stores the accumulated reciprocal and linear fair-area proxies separately from the raw model log-contract value. For both linear channels, the independent conditional Gaussian contribution uses the established Heston noise split:

$$\begin{aligned} \Delta v_{\text{independent}}^{\xi} &= (1 - \rho^2) \sigma \mathbb{E}_{\text{d}} \left[\int_{\text{step}} \xi_s ds \middle| \xi_{\text{start}} \right], \\ \Delta v_{\text{independent}}^{\Xi} &= (1 - P^2) \Sigma \mathbb{E}_{\text{d}} \left[\int_{\text{return step}} \Xi_s ds \middle| \mathcal{I} \right]. \end{aligned} \quad (4.350)$$

For the ordinary ξ channel, this restores the existing Heston independent share with $H = 2\sigma$ and the affine forward-mean area; its analytically normalized Gamma endpoint factor is retained. Subtracting endpoint entropy from the full linear area instead was an unrequested reciprocal residual prescription. Near unit correlation, intervals of order 10^{-16} – 10^{-12} exposed subtractive cancellation in the endpoint-mean and entropy calculation. Restoring the Heston share removes that residual subtraction, rather than clipping it or changing random draws or the time schedule.

For the Ξ return channel, the endpoint and first-return duration factors together carry the correlated contribution. For the exact continuous likelihood, the entropy chain rule separates its return-duration and conditional-path parts; their total expected log-contract contribution equals $P^2 \Sigma \mathbb{E}[\int_{\text{return}} \Xi_s ds]$. Subtracting only endpoint entropy from the full linear variance and then retaining the duration likelihood would allocate duration entropy again. The corrected independent share avoids that allocation.

The implemented ordinary and bridge endpoint factors are normalized Gamma mean-matching approximations, rather than the full continuous conditional-path likelihood. Their combined entropy need not equal the correlated budget exactly at finite N . Compare the raw log-contract value, including the duration

weight, with the stored fair-area proxy under refinement of the configured fixed step count; analytic forward normalization alone does not establish variance convergence. Both independent linear-channel shares now use (4.350); the state sampler, random draws, fixed full-cycle schedule and reciprocal frozen-sum formula remain unchanged. Scores obtained before the ordinary- ξ restoration are historical until repriced with this prescription. The initial market-pricing configuration has $\Sigma = 0$, so its duration factor is identically one. No final market-calibration or VIX consistency claim is made here.

In magic-time units, ξ, Ξ and τ have time dimension; κ, K have inverse-time dimension; σ, Σ have inverse-time-squared dimension; and $\aleph, \theta, \Theta, \rho, P$ are dimensionless. The normalized physical clock in (4.343) has time dimension; its normalization must be distinguished from the prescribed surrogate's mean magic-to-physical rate.

A reciprocal model with three clocks

The magic-time extension can be enlarged by allowing the reciprocal state and stock to respond directly to the observing driver's physical boundary clock. The separate implementation retains the earlier model as an exact special case; calibration of the additional channel is a separate task. Retain the magic-time clock τ , the natural rough clock $\mathfrak{r}_\tau^\xi$, and the mean-normalized physical clock C_τ from (4.243). The natural driver ξ and observing driver χ retain their existing dynamics. The Brownian innovation of χ remains independent of the natural driver and of the reciprocal-state and stock innovations. The observed system is still evaluated at $\tau(u(T))$, with the deterministic physical map applied before taking this inverse clock.

Use the superscript C to identify coefficients multiplying the *normalized* clock, and the subscript χ to identify its observing driver. The existing natural rough coefficients are unchanged. The domestic drift and covariance specification is

$$\begin{aligned}
\mathbb{E}_d[d\zeta] &= \left(\frac{1}{2} \mathfrak{p}_{\zeta, \tau} + \alpha - \beta \zeta \right) d\tau \\
&\quad + \left(\frac{1}{2} \mathfrak{p}_{\zeta, \mathfrak{r}} + \aleph - \beth \zeta \right) d\mathfrak{r}^\xi \\
&\quad + \left(\frac{1}{2} \mathfrak{p}_{\zeta, \chi}^C + \alpha_\chi^C - \beta_\chi^C \zeta \right) dC, \\
d\langle \zeta \rangle &= \zeta \left(\mathfrak{p}_{\zeta, \tau} d\tau + \mathfrak{p}_{\zeta, \mathfrak{r}} d\mathfrak{r}^\xi + \mathfrak{p}_{\zeta, \chi}^C dC \right), \\
\mathbb{E}_d[dS] &= 0, \\
\frac{d\langle S \rangle}{S^2} &= \left(\frac{\mathcal{H}\xi}{2} + \frac{1}{\zeta} \right) d\tau + \frac{\mathfrak{p}_S}{\zeta} d\mathfrak{r}^\xi + \frac{\mathfrak{p}_{S, \chi}^C}{\zeta} dC, \\
d\langle \zeta, S \rangle &= -S \left(\rho \sqrt{\mathfrak{p}_{\zeta, \tau}} d\tau + \mathfrak{n} \sqrt{\mathfrak{p}_{\zeta, \mathfrak{r}} \mathfrak{p}_S} d\mathfrak{r}^\xi + \rho_\chi \sqrt{\mathfrak{p}_{\zeta, \chi}^C \mathfrak{p}_{S, \chi}^C} dC \right).
\end{aligned} \tag{4.351}$$

The coefficient of $d\tau/\zeta$ in stock variance remains one, retaining the current scale choice for ζ . The optional Heston contribution and its natural-driver covariance are retained: $d\langle\xi, S\rangle = \rho_\xi S\sqrt{\mathcal{H}}\xi d\tau$. The cross-state covariations $d\langle\xi, \zeta\rangle$, $d\langle\chi, \zeta\rangle$, $d\langle\chi, S\rangle$, and $d\langle\xi, \chi\rangle$ are zero. Separate clock-channel innovations produce the stated covariance sums. The signs of ρ , $\mathfrak{N}$, and ρ_χ describe variance–stock correlations, opposite to reciprocal-state–stock correlations, as in (4.107).

The implementation names for the five new normalized-clock coefficients are `evar_S_physical`, `evar_zeta_physical`, `alpha_physical`, `beta_physical`, and `rho_physical`. These names distinguish them from the existing magic-time and natural rough parameters.

Positivity and the two numeraires. Diffusion strengths, stock-variance weights, and reversions are nonnegative. The magic-time state strength satisfies $\Psi_{\zeta,\tau} > 0$, and the correlations satisfy

$$|\rho| \leq 1, \quad |\mathfrak{N}| \leq 1, \quad |\rho_\chi| \leq 1.$$

Each clock-channel covariance matrix is then positive semidefinite. Retain the compulsory floor $\vartheta = 1$ in the one-half-diffusion terms of the state drifts; the actual Feller characteristics include their excess immigration. Under the stock numeraire these excesses become

$$\begin{aligned} \alpha_f &= \alpha - \rho\sqrt{\Psi_{\zeta,\tau}}, \\ \mathfrak{N}_f &= \mathfrak{N} - \mathfrak{N}\sqrt{\Psi_{\zeta,\tau}\Psi_S}, \\ \alpha_{\chi,f}^C &= \alpha_\chi^C - \rho_\chi\sqrt{\Psi_{\zeta,\chi}^C\Psi_{S,\chi}^C}. \end{aligned} \tag{4.352}$$

Nonnegative domestic and stock excesses give conservative channelwise nonattainment restrictions. In the new active state-diffusion channel, the domestic and stock characteristics are $1 + 2\alpha_\chi^C/\Psi_{\zeta,\chi}^C$ and $1 + 2\alpha_{\chi,f}^C/\Psi_{\zeta,\chi}^C$, respectively. Use strictly positive excesses in every active state-diffusion channel when requiring finite inverse-state expectations in the corresponding numeraire. The critical characteristic one is a nonattainment floor, not a finite-inverse-moment guarantee. When a diffusion strength is zero, apply the drift restrictions without dividing by that strength.

With the optional Heston component, the natural driver’s stock-numeraire reversion is $\kappa - \rho_\xi\sqrt{\mathcal{H}}$; its immigration θ is unchanged. The observing driver has no stock covariance, so its parameters and physical-clock normalization are unchanged by this numeraire change. The separate implementation reuses the current conditional stock kernel and respects that kernel’s nonpositive correlation restriction, although the covariance specification itself permits either sign subject to (4.352).

Canonical boundary coefficients and scaling. The normalized clock C has engineering units of time. Consequently $\Psi_{S,\chi}^C$, $\Psi_{\zeta,\chi}^C$, and α_χ^C are dimensionless,

whereas β_χ^C has units τ^{-1} . To express the same channel directly through the canonical boundary clock $\mathfrak{r}^\chi$, define instead

$$\begin{aligned}\mathfrak{v}_{S,\chi} &= \frac{\Gamma(\theta_\chi)}{(1-\theta_\chi)\kappa_\chi^{\theta_\chi}} \mathfrak{v}_{S,\chi}^C, \\ \mathfrak{v}_{\zeta,\chi} &= \frac{\Gamma(\theta_\chi)}{(1-\theta_\chi)\kappa_\chi^{\theta_\chi}} \mathfrak{v}_{\zeta,\chi}^C, \\ \mathfrak{N}_\chi &= \frac{\Gamma(\theta_\chi)}{(1-\theta_\chi)\kappa_\chi^{\theta_\chi}} \alpha_\chi^C, \\ \mathfrak{J}_\chi &= \frac{\Gamma(\theta_\chi)}{(1-\theta_\chi)\kappa_\chi^{\theta_\chi}} \beta_\chi^C.\end{aligned}\tag{4.353}$$

The canonical strengths and immigration have intrinsic powers τ^{θ_χ} , and $\mathfrak{J}_\chi$ has power $\tau^{\theta_\chi-1}$. These differ from the natural rough powers τ^θ and $\tau^{\theta-1}$ because the associated boundary driver is different. The normalized and canonical parameterizations are equivalent for $0 < \theta_\chi < 1$; their coefficient units must not be interchanged. Feller ratios are unchanged by this conversion. Short-horizon physical homogeneity remains distinct from engineering units and from increments started at a positive state, as discussed for the inverse observing clock. In observed physical time, $C_{\tau(u)} = u$, so the direct C channel runs continuously in u . In particular, its proportional reversion acts directly in physical time: when $\beta_\chi^C > 0$, its own attenuation duration is $1/\beta_\chi^C$. Applying the inverse-clock Laplace exponent to that rate would describe a different channel.

The relative clock scaling and the Hurst index. At their boundaries, before mean reversion introduces its characteristic scale, the natural and physical boundary clocks have magic-time scaling powers $1-\theta$ and $1-\theta_\chi$, respectively. Composing the natural clock with the inverse physical clock therefore gives the typical scaling

$$\mathfrak{r}_{\tau(u)}^\xi \sim u^{(1-\theta)/(1-\theta_\chi)}, \quad H_{\text{eff}} = \frac{1-\theta}{2(1-\theta_\chi)} = \frac{1}{2} \left(1 - \frac{\theta-\theta_\chi}{1-\theta_\chi} \right).\tag{4.354}$$

Here H_{eff} is the Hurst index suggested by an independent Brownian innovation in this rough clock. The relevant separation of the two clock powers is their ratio, or equivalently the normalized difference in the last expression, rather than the difference of the two Feller parameters alone. For example, keeping $H_{\text{eff}} = 0.1$ requires $\theta = 0.8 + 0.2\theta_\chi$. Increasing θ_χ with θ fixed increases this index; increasing both parameters according to that relation preserves it while changing the observing-clock jump law.

This statement concerns near-boundary distributional scaling. A positive initial ξ introduces a first-contact delay; mean reversion introduces crossovers, and the smooth clock channels can dominate observed volatility. The inverse physical clock creates jumps, so the complete observed process does not acquire the global continuous-path regularity of fractional Brownian motion with this Hurst index. Neither replacing the random horizon by its mean nor keeping a fixed unresolved-jump cutoff establishes this scaling. A nonlinear fitted map $u(T)$ also changes the relation between this physical-clock index and an index estimated in calendar time. When $\xi_0 = 0$ and $\theta > \theta_\chi$, the stable limiting magic horizon has the finite fractional moment needed for the expected natural rough clock to have the same leading physical-time power. At equality, the mean instead has a logarithmic correction; for $\theta < \theta_\chi$, rare finite magic excursions can give a leading mean proportional to u , even though the typical clock scale has the power in (4.354).

When changing θ in calibration, the units of the natural rough coefficients change. The separate scale helper retains their intrinsic reference durations as one explicit starting convention: for a coefficient with magic-time power a , retain $c^{1/a}$ and set the new coefficient to that duration raised to its new power. This preserves reference units, not prices or expected rough-clock increments; each proposed parameter pair must still be priced and compared with the same market objective.

A second starting convention instead retains nominal stationary clock-averaged coefficients in fixed numerical magic-time units. For $\kappa > 0$, the canonical natural rough clock has mean rate

$$\lim_{\tau \rightarrow \infty} \frac{\mathbb{E}[\mathfrak{r}_\tau^\xi]}{\tau} = \frac{(1 - \theta)\kappa^\theta}{\Gamma(\theta)}.$$

Multiply $\boldsymbol{\psi}_S$, $\boldsymbol{\mathfrak{N}}$, $\boldsymbol{\psi}_{\zeta, \mathfrak{r}}$, and $\boldsymbol{\mathfrak{I}}$ by the ratio of this rate at the old parameters to its rate at the new parameters, and set $\xi_{0, \text{new}} = \xi_{0, \text{old}} \kappa_{\text{old}} / \kappa_{\text{new}}$. This preserves their nominal averaged channel rates at frozen state, the rough Feller ratio, and the initial dimensionless location $\kappa \xi_0$. Optionally adjust $\mathcal{H}$ to retain $\mathcal{H}\theta/(2\kappa)$, the stationary magic-time mean of its Heston variance contribution. These are starting conventions, not invariances of the coupled state dynamics: state-clock dependence, actual moments, short-horizon shapes, and prices must be recomputed. The separate `three_clock_regime_starts.py` helper implements this alternative to retaining reference durations.

Nesting and simulation order. Setting all four new physical strengths and drift coefficients to zero recovers the current magic-time model, leaving every existing magic-time and natural rough parameter unchanged; ρ_χ is then irrelevant. The implementation explicitly delegates this case to the current simulator to retain its random inputs and prices exactly. At $\theta_\chi = 0$ define $C_\tau = \tau$ directly, rather than

evaluating the singular canonical normalizer. Nonzero physical coefficients then coalesce with the magic-time channel: diffusion strengths, stock weights, excess immigrations, reversions, and covariance coefficients add. Only those totals identify the coalesced channel, not its separate physical and magic-time amplitudes. Feller restrictions alone therefore do not remove this identification issue.

For positive θ_χ , C is constant throughout a skipped excursion of χ . The complete existing magic-time and natural rough system must evolve across its magic-time jump, with *zero* C increment during that excursion. The direct physical channel evolves between these jumps. Dividing an observation's physical increment by its sampled magic-time increment and smearing that ratio through the excursion would change the model and turn continuous physical-channel effects into jump effects.

A simulator could retain resolved observing-clock jump times and sizes, evolving the physical channel between those events. The separate implementation instead uses a refined physical grid and a symmetric splitting step: a physical half-step, the complete sampled magic-time and natural rough evolution, and a second physical half-step. This approximates the ordering of noncommuting channels and requires a physical-grid refinement check in addition to the existing latent-grid and unresolved-jump checks. It is not an exact replacement of the joint clock ordering. The implementation files are `three_clock_reciprocal.py` and `three_clock_reciprocal_market.py`; their regression checks include exact recovery of the earlier prices when the new channel is inactive. A coarse calibration improvement is not sufficient to replace an independently checked incumbent.

Cox-Ingersoll-Ross Model with Quasi-Rough Extension

The conventional CIR model can be formulated by making the instantaneous interest rate a linear function of the same affine process:

$$r = \frac{\nu^2}{2}\xi + r^{\min}.$$

Here, we choose the factor ν so it can be interpreted as a standard volatility of the square root process:

$$d\langle r \rangle = \nu^2(r - r^{\min})dt$$

As in Eq. 3.36, we will augment it with a rough component so that the pricing kernel becomes

$$\Phi_{t'} = \Phi_{t_0} \exp \left(\frac{\nu^2}{2} \int_{t_0}^{t'} \xi dt + r^{\min} \delta t + \mathfrak{r} \delta \mathfrak{t} \right). \quad (4.355)$$

The rough component does not affect the Kolmogorov equation outside the boundary, so that it remains the same as for the conventional CIR model:

$$\left(\partial_t + (\theta - \kappa\xi) \partial_\xi + \xi \partial_\xi^2 - \left(\frac{\nu^2}{2} \xi + r^{\min} \right) \cdot \right) \psi = 0. \quad (4.356)$$

Instead, the rough rates contribute via the boundary condition Eq. 3.38. Alternatively, one can use the dual Fokker-Planck equation:

$$\left(-\partial_t + \partial_\xi(\kappa\xi - \theta) \cdot + \partial_\xi^2 \xi \cdot - \left(\frac{\nu^2}{2} \xi + r^{\min} \right) \cdot \right) \psi = 0 \quad (4.357)$$

with the boundary condition Eq. 3.40 for Arrow-Debreu prices. The rough component allows us to have large long rates in combinations with low or even zero short rates. The latter occurs if we remove the smooth component by setting $\nu = r^{\min} = 0$. This feature helps to model the low-rate regime. In an application to credit, this would produce a class of models resembling the Merton capital-structure model. In the edge case $\mathfrak{T} = \infty$, the default occurs immediately after the boundary is reached, increasing the similarity to the conventional Merton model.

Jump diffusion

The methodology can also be applied to model jumps. For this purpose, we should consider a process of type Eq. 4.40, Eq. 4.59, or of a more generic form, while replacing physical time with an abstract clock τ , as proposed in [7]:

$$d\xi_\tau = (\theta - \kappa\xi_\tau) d\tau + \sqrt{2\xi_\tau} dW_\tau. \quad (4.358)$$

Then we can model the physical time as the edge time for this process.

$$t_\tau = \mathfrak{r}_\tau.$$

As we have demonstrated, this would make the clock τ a pure jump process in physical time. Consequently, other processes, such as asset prices, rates, etc., defined in terms of the clock τ would also incur jumps. The distributions of this model can be obtained via the inverse Laplace transform sourced from Eq. 4.53 with complex arguments. Alternatively, we can solve the model by applying the Kolmogorov equation derived in the clock τ . The abstract clock τ cannot directly enter the pricing problem. Therefore, the relevant solutions must be stationary in terms of τ , which eliminates the time derivative term from the Kolmogorov equation, effectively making it elliptic instead of diffusive. The relation to physical time comes exclusively from the boundary condition

$$\lim_{\xi \rightarrow 0} \partial_t \psi = {}^1\psi.$$

In addition, the process would only be observable at the boundary. This methodology would allow us to handle jump diffusion without integro-differential equations.

Hybrid multi-clock models

The methodology can be used to simulate various clocks that naturally occur in financial models, such as variance clocks, money market accrual clocks, default probability accrual clocks, etc. When modeling these clocks, we can incorporate rough components in addition to the traditional smooth components, representable as integrals over continuous positive processes. Such rough components can be modeled as a linear combination of edge times obtained from a set of affine processes. In addition, the physical clock can be modeled in the same fashion, which would allow one to incorporate jumps.

Summary

- We extend conventional Brownian motions by incorporating additional clocks associated with the local time spent at the process boundary.
- We speculate that the methodology can emulate a rough diffusion. It also provides a computationally efficient way to introduce jumps.
- We demonstrate that the approach is capable of handling many stylized facts pertaining to stochastic volatility.

A Joint distribution of ξ_t and $\mathfrak{r}_t$ from the boundary

We consider the mean-reversion-free affine process Eq. 4.1 started at the boundary, $\xi_0 = 0$. This appendix collects the joint law of the terminal state ξ_t and the edge time $\mathfrak{r}_t$ in the form of a one-dimensional integral, obtained by conditioning on the last hit of the boundary.

Last hit of the boundary

Introduce the auxiliary variable

$$u_t \equiv \sup \{s \leq t : \xi_s = 0\}, \quad 0 \leq u_t \leq t. \quad (\text{A.1})$$

By the Dynkin–Lamperti generalised arcsine law [11] for the $(1 - \theta)$ -stable subordinator underlying $\mathfrak{r}$ (see Eq. 4.12), the ratio u_t/t is Beta($1 - \theta, \theta$), with density

$$f_{u_t}(u) = \frac{\sin(\pi\theta)}{\pi} u^{-\theta} (t - u)^{\theta-1}, \quad 0 < u < t. \quad (\text{A.2})$$

At $\theta = 1/2$, Eq. A.2 collapses to the classical arcsine density $1/[\pi\sqrt{u(t-u)}]$ of the last zero of a Brownian motion on $(0, t)$, consistent with the identification $\xi = (\eta - W)^2/2$ available at $\theta = 1/2$.

Conditional independence given u_t

The strong Markov property at u_t implies that the excursion of ξ on $[u_t, t]$ (which determines ξ_t) is independent of the trajectory on $[0, u_t]$ (which determines $\mathfrak{r}_t = \mathfrak{r}_{u_t}$). Consequently,

$$\xi_t \perp\!\!\!\perp \mathfrak{r}_t \mid u_t,$$

and the joint conditional density factorises as

$$f_{\xi_t, \mathfrak{r}_t | u_t}(\xi, \mathfrak{r} \mid u) = f_{\xi_t | u_t}(\xi \mid u) f_{\mathfrak{r}_t | u_t}(\mathfrak{r} \mid u). \quad (\text{A.3})$$

Conditional distribution of ξ_t

Given $u_t = u$, the process ξ on $(u, t]$ is the excursion of the affine of index θ from an entrance at 0, conditioned on non-return within $(u, t]$. Its density at excursion age $s = t - u$ is fixed by the requirement that the last-hit decomposition

$${}^0G(t, \xi \mid 0) = \int_0^t f_{u_t}(u) f_{\xi_t | u_t}(\xi \mid u) du$$

reproduce the entrance density Eq. 4.15. Matching Laplace transforms in ξ (both sides depend on ξ only through the argument λ) leads to a Gauss hypergeometric identity, ${}_2F_1(1, \theta; 1; -\lambda t) = (1 + \lambda t)^{-\theta}$, that pins down the shape parameter of the meander to unity, independently of θ :

$$f_{\xi_t | u_t}(\xi \mid u) = \frac{e^{-\xi/(t-u)}}{t - u}, \quad \xi > 0. \quad (\text{A.4})$$

That is, $\xi_t \mid u_t = u \sim \text{Exp}(\text{rate} = 1/(t - u))$, of mean $t - u$. Independence of the shape from θ is a coincidence of the arcsine-and-entrance-density pair, not a θ -uniform statement about the ambient excursion measure.

Conditional distribution of $\mathfrak{r}_t$

Since $\mathfrak{r}_t = \mathfrak{r}_{u_t}$, the conditional law of $\mathfrak{r}_t$ given $u_t = u$ is the density of the physical time $t_{\mathfrak{r}}$ at u (a $(1 - \theta)$ -stable subordinator by Eq. 3.5, Eq. 4.12), normalised as a function of $\mathfrak{r}$ for fixed u :

$$f_{\mathfrak{r}_t|u_t}(\mathfrak{r} | u) = \Gamma(1 - \theta) u^\theta f_{t_{\mathfrak{r}}}(u), \quad \mathfrak{r} > 0. \quad (\text{A.5})$$

Taking the first moment via the Mellin transform of the one-sided $(1 - \theta)$ -stable law [13] yields the conditional expectation in closed form:

$$\mathbb{E}[\mathfrak{r}_t | u_t = u] = \frac{\Gamma(1 - \theta) \Gamma(2 - \theta)}{\Gamma(\theta) \Gamma(2 - 2\theta)} u^{1-\theta}. \quad (\text{A.6})$$

Integrating Eq. A.6 against the arcsine density Eq. A.2 reproduces the marginal expectation Eq. 4.27 at $\xi_0 = 0$, $\mathbb{E}[\mathfrak{r}_t | \xi_0 = 0] = t^{1-\theta}/\Gamma(\theta)$.

At $\theta = 1/2$, Eq. A.5 reduces to a Rayleigh law with parameter $\sqrt{u/2}$,

$$f_{\mathfrak{r}_t|u_t}(\mathfrak{r} | u) = \frac{2\mathfrak{r}}{u} \exp\left(-\frac{\mathfrak{r}^2}{u}\right),$$

and Eq. A.6 gives $\mathbb{E}[\mathfrak{r}_t | u_t = u] = \frac{1}{2}\sqrt{\pi u}$, matching the closed-form joint density obtained via the Brownian representation of Eq. 2.1.

Joint density

Combining Eq. A.2, Eq. A.3, Eq. A.4, and Eq. A.5,

$$f_{\xi_t, \mathfrak{r}_t}(\xi, \mathfrak{r} | \xi_0 = 0) = \int_0^t f_{t_{\mathfrak{r}}}(u) \frac{(t - u)^{-(1-\theta)}}{\Gamma(\theta)} \frac{e^{-\xi/(t-u)}}{t - u} du. \quad (\text{A.7})$$

The three factors in the integrand are, respectively: the density of the physical time $t_{\mathfrak{r}}$ at u ; the tail Lévy measure of $t_{\mathfrak{r}}$ at $t - u$ (probability rate of an excursion of length exceeding $t - u$); and the meander density of ξ_t at excursion age $t - u$. Except at $\theta = 1/2$, where $t_{\mathfrak{r}}$ has the Lévy distribution and Eq. A.7 reduces to elementary form, the density of $t_{\mathfrak{r}}$ is a scaled one-sided $(1 - \theta)$ -stable density [13] and the integral does not close in elementary functions.

B Terminal asset law at a specified next touch

This appendix records the endpoint-first construction for the smooth joint (ξ, S) process. Let T be the specified next touch time, so that $\xi_T = 0$, and define the integrated affine area

$$A_T \equiv \int_0^T \xi_s ds. \quad (\text{B.1})$$

For the smooth Heston component of Eq. 4.62, substitution of

$$dW_t^\xi = \frac{d\xi_t - (\theta - \kappa\xi_t)dt}{\sqrt{2\xi_t}}$$

and integration to the touch give

$$\begin{aligned} \log \frac{S_T}{S_0} = & -\frac{\rho\sqrt{\mathcal{H}}}{2}(\xi_0 + \theta T) + \left(\frac{\rho\sqrt{\mathcal{H}}}{2}\kappa - \frac{\mathcal{H}}{4} \right) A_T \\ & + \sqrt{\frac{(1-\rho^2)\mathcal{H}}{2}} \int_0^T \sqrt{\xi_s} dW_s^\perp. \end{aligned} \quad (\text{B.2})$$

Consequently, conditional on the ξ path and hence on A_T , the terminal log price is Gaussian with mean

$$m_T(A) = -\frac{\rho\sqrt{\mathcal{H}}}{2}(\xi_0 + \theta T) + \left(\frac{\rho\sqrt{\mathcal{H}}}{2}\kappa - \frac{\mathcal{H}}{4} \right) A \quad (\text{B.3})$$

and variance

$$v_T(A) = \frac{(1-\rho^2)\mathcal{H}}{2} A. \quad (\text{B.4})$$

The conditional terminal distribution at the specified next touch is thus a normal variance-mean mixture, determined by the conditional area law

$$\mathcal{L}_T(q) \equiv \mathbb{E}[e^{-qA_T} | \xi_0, T_{\text{touch}} = T]. \quad (\text{B.5})$$

For example, its characteristic function is

$$\mathbb{E}\left[e^{iu \log(S_T/S_0)} | T_{\text{touch}} = T\right] = e^{-iu\rho\sqrt{\mathcal{H}}(\xi_0 + \theta T)/2} \mathcal{L}_T\left(\frac{u^2(1-\rho^2)\mathcal{H}}{4} + iu\left(\frac{\mathcal{H}}{4} - \frac{\rho\sqrt{\mathcal{H}}}{2}\kappa\right)\right). \quad (\text{B.6})$$

The transform Eq. B.5 is obtained as the ratio of the area-killed first-touch density to the ordinary first-touch density, i.e. by a killed Feynman–Kac boundary-flux calculation. Let $x = \log S$ and let

$$h_{T,y}(t, \xi, x) \equiv f_{T_{\text{touch}}, X_T}(T, y | t, \xi, x) \quad (\text{B.7})$$

be the joint terminal touch/log-price density, obtained by Fourier inversion of Eq. B.6 multiplied by the first-touch density. Conditioning on the generated pair (T, y) is the Doob transform with this random Radon–Nikodym derivative. The covariance matrix of (ξ, x) is

$$a(\xi) = \xi \begin{pmatrix} 2 & \rho\sqrt{\mathcal{H}} \\ \rho\sqrt{\mathcal{H}} & \mathcal{H}/2 \end{pmatrix},$$

so its conditional drifts are $b^h = b + a\nabla \log h_{T,y}$. In particular,

$$\begin{aligned}\mathbb{E}^h[d\xi] &= \left(\theta - \kappa\xi + 2\xi\partial_\xi \log h_{T,y} + \rho\sqrt{\mathcal{H}}\xi\partial_x \log h_{T,y} \right) dt, \\ \frac{\mathbb{E}^h[dS]}{S} &= \left(\rho\sqrt{\mathcal{H}}\xi\partial_\xi \log h_{T,y} + \frac{\mathcal{H}\xi}{2}\partial_x \log h_{T,y} \right) dt.\end{aligned}\tag{B.8}$$

This supplies the endpoint-first construction: generate T , generate S_T from Eq. B.6, then generate the joint (ξ, S) bridge using Eq. B.8. The implementation provides the terminal transform, its Fourier density, and the numerical Doob drift; replacing the legacy stepwise asset transition with the full bridge sampler remains pending.

C Implementation with Pytorch

```
class rough_stoch_vol_model:
    def __init__(self, **kwargs):
        self.aleph = kwargs.get('aleph')
        if self.aleph is None:
            self.rough_time = kwargs.get('rough_time', 1/12)
            self.aleph = sqrt(self.rough_time)
        else:
            self.rough_time = self.aleph ** 2

        self.xi = torch.tensor(self.rough_time * kwargs.get('init_state', 0.5))
        self.rough_var = kwargs.get('rough_vol', 0.2)
        elasticity = kwargs.get('elasticity')

        if elasticity is None:
            self.alpha = kwargs.get('alpha', 1)
        else:
            self.alpha = 1/(1 - elasticity)

        self.beta = kwargs.get('beta', 0)
        self.gamma = kwargs.get('gamma', 1)
        self.time_index = 0
        self.horizon = kwargs.get('horizon', 3)
        self.num_steps = kwargs.get('num_steps', 2**9)
        self.num_paths = kwargs.get('num_paths', 2**19)
        self.antithetic = kwargs.get('antithetic', True)
        if self.antithetic: self.num_paths >>= 1
        self.time_step = self.horizon/self.num_steps
```

```

self.step_aleph = self.aleph * sqrt(self.time_step)
seed = kvargs.get('seed')
if not seed is None: torch.manual_seed(seed)
self.rn = generate_random(2, self.num_steps, self.num_paths)
self.correlation = kvargs.get('correlation', 0)
if self.antithetic: self.rn = make_antithetic(self.rn)
self.adjustment = torch.tensor(sqrt(log(2)/2) * self.step_aleph)
self.time_index = 0
self.x = 0

def smooth_vol(self):
    return (self.gamma/self.rough_time * self.xi).pow(self.alpha) * (
        self.beta/self.rough_time * self.xi).exp()

def propagate(self,T):
    time_index = min(int(T/self.time_step), self.rn.size()[0])

    for r in self.rn[self.time_index:time_index]:
        smooth_vol = self.smooth_vol() * sqrt(self.time_step)
        smooth_var = smooth_vol.pow(2)
        xi = self.xi - self.time_step + r[0] * self.step_aleph
        self.xi = xi.max(self.adjustment)
        total_var = (self.xi - xi) * self.rough_var + smooth_var
        isync_var = total_var - smooth_var * self.correlation ** 2
        self.x = self.x + self.correlation * smooth_vol * r[0] \
            + isync_var.sqrt() * r[1] - 0.5* total_var

    self.time_index = time_index

def asset_price(self):
    return self.x.exp()

```

D Discretization of the coupled ξ – S dynamics under a touch-time ξ simulator

We simulate ξ via the touch-time / meander construction of subsection 4.1. On each step $[t_i, t_{i+1}]$ the touch-time simulator returns the sampled ξ_{i+1} together with a flag identifying whether the current excursion continues through the step (*case B*, $T_i > t_{i+1}$) or a touch of the boundary occurred inside the step (*case A*, $T_i \leq t_{i+1}$,

where T_i is the pending first-touch time entering the step).

We then evolve S conditional on the sampled ξ trajectory using the SDE-consistent Ito integration of the classic-form dynamics Eq. 4.62. Substituting $dW^\xi = (d\xi - (\theta - \kappa\xi) dt)/\sqrt{2\xi}$ into $dS/S = \sqrt{\mathcal{H}\xi/2}(\rho dW^\xi + \sqrt{1 - \rho^2} dW^\perp)$ gives the closed-form conditional log- S dynamics

$$d \log S = -\frac{1}{2} \frac{\mathcal{H}\xi}{2} dt + \frac{\rho\sqrt{\mathcal{H}}}{2} (d\xi - (\theta - \kappa\xi) dt) + \sqrt{(1 - \rho^2) \frac{\mathcal{H}\xi}{2}} dW_t^\perp, \quad (\text{D.1})$$

where the CIR diffusion coefficient $\sqrt{2\xi}$ cancels the $\sqrt{\mathcal{H}\xi/2}$ prefactor and leaves the genuinely constant $\rho\sqrt{\mathcal{H}}/2$. This is the joint (ξ, S) analogue of Eq. 4.39: the Girsanov drift acquired by $\log S$ under conditioning on ξ is the deterministic scalar $\rho\sqrt{\mathcal{H}}(1 - \theta) dt$ in the never-touch limit, and reduces to Eq. D.1 pathwise.

One-step discretization. For each step,

$$\Delta \log S = [\Delta \log S]_{\mathcal{H}} + \Delta \log S^{\text{other}}, \quad (\text{D.2})$$

where the $\mathcal{H}$ -channel contribution is

$$[\Delta \log S]_{\mathcal{H}} = -\frac{1}{2} \frac{\mathcal{H}\xi_i}{2} \delta t + \frac{\rho\sqrt{\mathcal{H}}}{2} (\xi_{i+1} - \xi_i - (\theta - \kappa\xi_i) \delta t) + \sqrt{(1 - \rho^2) \frac{\mathcal{H}\xi_i}{2}} \delta t z_H^\perp, \quad (\text{D.3})$$

with $z_H^\perp \sim \mathcal{N}(0, 1)$ independent across steps and paths, and $\Delta \log S^{\text{other}}$ collects the smooth- B contribution $-\frac{1}{2} B_i \delta t + \sqrt{B_i \delta t} z_B$ and the rough S contribution $-\frac{1}{2} \boldsymbol{\Psi}_S \delta \mathbf{r} + \sqrt{\boldsymbol{\Psi}_S \delta \mathbf{r}} z_S^{\text{rough}}$ (with z_S^{rough} optionally correlated with z_B^{rough} through $\mathfrak{D}$).

Touching steps. On a step where the excursion completes ($T_i \leq t_{i+1}$, "case A"), the process crosses the boundary $\xi = 0$. Although the diffusion coefficient $\sqrt{\mathcal{H}\xi/2}$ formally vanishes at $\xi = 0$, this only makes the integrated variance $\int_{t_i}^{t_{i+1}} \mathcal{H}\xi_s/2 ds$ a sub-interval quantity, not zero: $\xi_i > 0$ at the start of the step, so the left-endpoint Ito compensator $-\frac{1}{2}(\mathcal{H}\xi_i/2)\delta t$ is $\mathcal{O}(\delta t)$, not $\mathcal{O}(\delta t^2)$. Dropping the $\mathcal{H}$ -channel contribution on case-A steps therefore removes an $\mathcal{O}(\delta t)$ compensator that does not cancel, introducing an $\mathcal{O}(1)$ per-step martingale bias proportional to the fraction of case-A steps. We consequently apply Eq. D.3 on *every* step regardless of case A/B; the ρ -driven $(\xi_{i+1} - \xi_i - \text{CIR drift} \cdot \delta t)$ term uses the touch-time sampled $\Delta\xi$ (which correctly averages to the CIR drift over case-A + case-B combined) so the correlated shock is unbiased.

Antithetic S-side sampling. The $z_H^\perp$, z_B , z_S^{rough} shocks are drawn as antithetic pairs: for N paths (even), $N/2$ Gaussians are drawn and the remaining $N/2$ are their sign-flipped copies. This halves the S-side variance of every function of $\log S$ that is smooth in the shocks, at no additional cost. The ξ path is shared across the pair (the touch-time Poisson-Gamma-Beta chain is not antitheticised).

Earlier martingale diagnostics. At the stage-1 fit of a EURUSD 7d smile, $\theta = 0.8$, $\kappa = 0$, $\xi_0 = 0.100$, $\mathcal{H} = 0.0456$, using $N = 2^{20} = 1,048,576$ antithetic paths and daily steps:

ρ	T	steps	$\mathbb{E}[S_T] - 1$	bias/SE	CV log-MGF diff ($\alpha = 0.5$)	implied IV shift
0	7d	7	$+1.7 \times 10^{-6}$	+0.25	$+4.9 \times 10^{-8}$	< 0.05 bps
0	30d	30	$+3.0 \times 10^{-6}$	+0.19	$+3.0 \times 10^{-7}$	< 0.15 bps
0	30d	120	-1.3×10^{-5}	-0.87	$+2.9 \times 10^{-8}$	< 0.05 bps
-0.5	30d	30	-3.4×10^{-6}	-0.22	$+2.6 \times 10^{-7}$	< 0.15 bps
-0.5	30d	120	-1.4×10^{-5}	-0.89	-1.4×10^{-8}	< 0.05 bps
-0.9	30d	30	-9.1×10^{-6}	-0.60	$+1.7 \times 10^{-7}$	< 0.15 bps
-0.9	30d	120	-9.8×10^{-6}	-0.64	$+2.0 \times 10^{-8}$	< 0.05 bps

These earlier diagnostics were obtained before the current conditional joint transition review. Subsequent experiments report a stable martingale bias that does not disappear under grid refinement. The table is retained as historical output only and must not be read as validation of the present stepwise joint (ξ, S) construction.

Verification of the ξ marginal. The touch-time simulator’s one-step law is verified against the closed-form CIR marginal $\xi_t \sim (t/2) \cdot \text{NC}\chi^2(2\theta, 2\xi_0/t)$ (at $\kappa = 0$) by direct comparison of first two moments and the Laplace transform $\mathbb{E}[\exp(-\lambda\xi_t)]$ at multiple λ . At $N = 2^{22}$ paths and $T = 7\text{d}/30\text{d}/90\text{d}$: first two moments match to $\lesssim 0.1\%$ (MC noise), Laplace transform to $\lesssim 3 \times 10^{-4}$ absolute across $\lambda \in \{1, 5, 20, 100\}$. The residual bias in the earlier version of the simulator, of order -4% in $\mathbb{E}[\xi_t]$ for $\xi_0 > 0$, is traced to a catastrophic-cancellation in the case- B chi-squared sampling of the form $1/s_i - 1/s_{\text{prev}}$ when the drawn touch time T is large; rewriting the same expression as $(s_{\text{prev}} - s_i)/(s_{\text{prev}} s_i) = (\bar{t}_i - \bar{t}_{\text{prev}})/(s_{\text{prev}} s_i)$ removes the underflow without changing the algorithm at exact arithmetic.

Higher-order refinements. The Ito integration in Eq. D.3 is exact conditional on the ξ path taken as a Gaussian approximation of $\Delta\xi$; the residual bias comes from the higher cumulants of the touch-time $\Delta\xi$ distribution. These vanish as $\delta t \rightarrow 0$ (confirmed by the $\mathcal{O}(\delta t)$ decay of the CV log-MGF diff in the table above). If higher precision is required, the CIR transition MGF has closed form and can be used to replace the Euler compensator $-\frac{1}{2}(\mathcal{H}\xi_i/2)\delta t$ with the exact per-step $-\log \mathbb{E}[\exp((\rho\sqrt{\mathcal{H}}/2)\Delta\xi) \mid \xi_i]$, enforcing per-step martingality exactly.

E Implementation-specific sampled-touch streaming method

This appendix records the implementation-specific method currently used by the sampled-touch pseudo-rough simulator. It is deliberately separate from the general derivations above. The affine process, boundary-touch construction, and conditional laws are inherited from the preceding sections; this appendix describes the particular numerical method and its practical corrections.

E.1 Method flavour

The implemented method is an *endpoint-conditioned, sampled-touch, streaming Markov method*. The physical-time affine state is advanced through the sampled next-touch construction. Once the next endpoint ξ_{i+1} has been generated, the asset is advanced conditional on the realised affine endpoint. The asset is not simulated on a separate unconditional touch-time path and is not conditioned on the touch time beyond its effect on the realised ξ state and the area approximation.

The method therefore has the following character:

- **Sampled touch structure:** pending touch times determine whether a step is a continuing excursion or contains a completed excursion.
- **Exact affine endpoint laws within the construction:** continuing excursions use the conditional noncentral chi-square law; completed excursions use the Beta last-touch law and Gamma(1) meander law.
- **Conditional asset update:** the correlated Heston increment is reconstructed from the sampled $\Delta\xi$, while an independent Gaussian supplies the residual Heston variance.
- **Streaming Monte Carlo:** only the current ξ , edge time, B , and $\log S$ are retained. Payoffs are evaluated at their expiry steps.

E.2 Sampled ξ transition and vectorised Gamma construction

For a continuing excursion, the noncentral chi-square transition is sampled through the Poisson–Gamma mixture

$$N \sim \text{Poisson}(\lambda/2), \quad G \sim \text{Gamma}(\nu/2 + N, 1), \quad \chi_\nu'^2(\lambda) = 2G. \quad (\text{E.1})$$

For a completed excursion, the terminal meander is $G_1 \sim \text{Gamma}(1, 1)$ multiplied by its conditional age scale. In the streaming implementation the two populations are assembled into one vector of Gamma shapes and scales, followed by one broadcast Gamma draw. This preserves the two distributions while avoiding separate touch and continuation sampling calls for a mixed batch.

E.3 Conditional streaming update of S

The smooth Heston channel is updated after the affine endpoint has been sampled. With $Z_i^\perp \sim N(0, 1)$,

$$\Delta \log S_i^H = \log R_i - \frac{1}{2} V_i^\perp + \sqrt{V_i^\perp} Z_i^\perp, \quad (\text{E.2})$$

where $\log R_i$ is the endpoint affine martingale correction and

$$V_i^\perp = \frac{(1 - \rho^2) \mathcal{H}}{2} A_{\xi, i}. \quad (\text{E.3})$$

The correlated part is contained in $\log R_i$ and uses the realised ξ_{i+1} endpoint. The independent residual uses the forward conditional mean area $A_{\xi, i}$ rather than a stored path integral.

For no touch, the forward projection starts at ξ_i . For a touched step, the contribution before the last touch is discarded and the projection starts from zero at t_{last} . For $\kappa = 0$,

$$A_{\xi, i} = \xi_i d + \frac{1}{2} \theta d^2 \quad (\text{no touch}), \quad A_{\xi, i} = \frac{1}{2} \theta d^2 \quad (\text{touch}), \quad (\text{E.4})$$

where $d = \Delta t_i$ in the first case and $d = t_{i+1} - t_{\text{last}}$ in the second. For $\kappa \neq 0$, with $\bar{\xi} = \theta / \kappa$,

$$A_{\xi, i} = \bar{\xi} d + (\xi_i - \bar{\xi}) \frac{1 - e^{-\kappa d}}{\kappa}, \quad (\text{no touch}), \quad (\text{E.5})$$

$$A_{\xi, i} = \bar{\xi} d - \bar{\xi} \frac{1 - e^{-\kappa d}}{\kappa}, \quad (\text{touch}). \quad (\text{E.6})$$

The $\kappa \rightarrow 0$ expressions are used directly to avoid cancellation.

E.4 Complete and extreme B treatments

The complete method adds a smooth physical-time B variance, a rough stock variance, and rough-channel correlation:

$$V_i = V_i^\perp + \int_{t_i}^{t_{i+1}} B_s ds + \psi_S \Delta \mathfrak{r}_i. \quad (\text{E.7})$$

Before the first touch and after the last touch, B follows its deterministic physical-time mean-reversion propagator. At a completed excursion the rough edge-time terminal value B_{last} is sampled. The cumulative conditional mean-reversion contribution over edge time is then converted to a uniform physical-time level on $[t_{\text{first}}, t_{\text{last}}]$, preserving the integrated expected contribution over that interval. The post-touch physical-time propagation starts from B_{last} .

The zero-rate limits are explicit. If $\beta_B = 0$,

$$B(t) = B_0 + \alpha_B t, \quad \int_0^d B(t) dt = B_0 d + \frac{1}{2} \alpha_B d^2. \quad (\text{E.8})$$

The same continuous limit is used for the rough edge-time rate $\mathfrak{z}_B = 0$. The extreme method is a separate approximation: it resets B to $\mathfrak{y}_B$ over the rough interval and does not include rough B diffusion or rough B - S correlation.

E.5 Direct- ξ elapsed-edge-time approximation

The alternative fast method does not use a touch-time bridge or the Bessel-function conditional touch probability. It simulates the affine endpoint directly and estimates the conditional expected edge-time increment from the two endpoint values. For $dt = t_i - t_0$ and small ξ_0/dt , the approximation used is

$$\begin{aligned} \mathbb{E}[\Delta \mathfrak{r} \mid \xi_{t_0} = \xi_0, \xi_{t_i} = \xi_i] &\approx \frac{\Gamma(1-\theta)\Gamma(2-\theta)}{\Gamma(\theta)\Gamma(2-2\theta)} dt^{1-\theta} \\ &\quad - \frac{\Gamma(2-\theta)}{\Gamma(2-2\theta)} \xi_0^{1-\theta} \\ &\quad - \xi_i^{1-\theta} \left(1 - \frac{\xi_0^{1-\theta}}{\Gamma(2-\theta)dt^{1-\theta}} \right). \end{aligned} \quad (\text{E.9})$$

It is useful to read this as four terms: a positive $dt^{1-\theta}$ contribution, a negative initial-endpoint correction, a negative terminal-endpoint correction, and a positive $\xi_0^{1-\theta}\xi_i^{1-\theta}$ cross-term divided by $dt^{1-\theta}$.

The derivation starts from the boundary-start conditional result,

$$\mathbb{E}[\Delta \mathfrak{r} \mid \xi_{t_0} = 0, \xi_{t_i} = \xi_i] \approx \frac{\Gamma(1-\theta)\Gamma(2-\theta)}{\Gamma(\theta)\Gamma(2-2\theta)} dt^{1-\theta} - \xi_i^{1-\theta}. \quad (\text{E.10})$$

For nonzero but small ξ_0 , the first-touch density is integrated against the post-touch boundary-start expression. Expanding that integral in the small parameter ξ_0/dt gives the initial correction $-\Gamma(2-\theta)\xi_0^{1-\theta}/\Gamma(2-2\theta)$ and the endpoint cross-term in Eq. E.9. This is an endpoint approximation for the fast direct- ξ method, not the exact bridge conditional expectation. The computed increment is floored at zero.

E.6 Streaming, antithetic shocks, and option evaluation

The simulation keeps one current state per path. The B and residual Heston shocks use antithetic sign patterns, while the touch-time affine chain is sampled from its own Gamma, Poisson, and Beta draws. Options sharing an expiry are evaluated together when the corresponding grid step is reached. Thus all strikes and call/put signs at an expiry use the same simulated path set and the same forward-mean estimate.

E.7 Implementation validation

The focused implementation suite contains tests for the touch and continuation Gamma laws, $\kappa = 0$ and nonzero- κ forward areas, the zero-rate B limits, cumulative rough- B contribution, and finite pure, complete, and extreme streaming prices. The current suite contains 13 passing tests.

For a pure-Heston benchmark with $\theta = 0.8$, $\kappa = 1$, $\xi_0 = 0$, $\mathcal{H} = 0.05$, 2^{20} shared paths, and 256 steps, the one-year results were:

ρ	$\mathbb{E}[S_1]$	$\max_{\lambda=0.1,\dots,0.9} \Sigma_{\text{MC}} - \Sigma_{\text{analytic}} $
0.0	0.99999996	6.73×10^{-6}
-0.5	0.99998810	4.22×10^{-6}
+0.5	1.00000820	5.81×10^{-6}
+0.9	1.00001054	3.71×10^{-7}
-0.9	0.99997733	1.39×10^{-6}

These pure-Heston tests are independent of the complete-model B channel.

F Sampled-touch simulation of the rough-asset reciprocal model

This appendix specifies the original rough-asset specialization of (4.64): $c_c = 0$ and $\Psi_S = 1$ in the original numerical time and reciprocal-state units. All reciprocal return tilts, residual variances, and Feller shifts in this appendix use that specialization. The calendar-time stock scheme is given separately in Appendix G. The implementation is separate from the lognormal- B simulator, with three modules in `py/`:

- `reciprocal_clock.py`: copied sampled-touch clock machinery;
- `sample_touch_reciprocal.py`: reciprocal state and asset updates;

- `rough_reciprocal_market_calibration.py`: pricing and calibration.

There is no B state in this scheme. The touch-time, ξ , and $\mathfrak{r}$ constructions are those of Appendix E; their probabilistic derivations are not repeated here.

The simulation order on each calendar-time step is: sample the touch information and ξ endpoint, propagate ζ using that clock information, then update S through two analytically normalized conditional-return factors and independent lognormal residuals. The central requirement is that the *discrete asset update* preserve the martingale condition without rescaling a realised Monte Carlo sample.

F.1 Clock information and retained state

At t_i a base path retains ξ_i , ζ_i , $\mathfrak{r}_i$, the pending next-touch time, and $\log S_i$. On $[t_i, t_{i+1}]$ the clock sampler returns ξ_{i+1} , $\Delta \mathfrak{r}_i$, and either a continuing-excursion flag or the first and last sampled touches, t_{first} and t_{last} . Continuing excursions retain their pending first touch and use the conditional noncentral Gamma law. A touched step uses the Beta last-touch law and the Gamma(1) terminal meander law, then samples a new pending touch from its new endpoint. This preserves the existing clock sampler, including its random-draw ordering.

The retained elapsed-edge-time approximation is

$$\Delta \mathfrak{r}_i = \frac{\Gamma(1-\theta)\Gamma(2-\theta)}{\Gamma(\theta)\Gamma(2-2\theta)} (\bar{t}_{\text{last}} - \bar{t}_{\text{first}})^{1-\theta} \quad \text{on a touched step,} \quad (\text{F.1})$$

and zero on a continuing excursion. Here differences in the transformed clock are calendar-time differences for $\kappa = 0$, and

$$\bar{t}_{\text{last}} - \bar{t}_{\text{first}} = \frac{e^{\kappa t_{\text{last}}} - e^{\kappa t_{\text{first}}}}{\kappa} \quad (\kappa \neq 0). \quad (\text{F.2})$$

Equation (F.1) is a numerical approximation of the elapsed edge time; it is not a draw of its full conditional bridge distribution.

F.2 Calendar drift outside the touched interval

When the rough clock is flat, ζ follows the deterministic flow

$$\mathcal{D}_d(\zeta) = e^{-\beta d} \zeta + \alpha \frac{1 - e^{-\beta d}}{\beta}, \quad \mathcal{D}_d(\zeta) = \zeta + \alpha d \quad (\beta = 0). \quad (\text{F.3})$$

Thus, on a no-touch step, $\zeta_{i+1} = \mathcal{D}_{\Delta t_i}(\zeta_i)$ and the reciprocal channel contributes neither stock variance nor a conditional-return correction. The Heston channel still operates on this step.

On a touched step the pre-touch state is

$$\zeta_{\text{first}} = \mathcal{D}_{t_{\text{first}}-t_i}(\zeta_i). \quad (\text{F.4})$$

The span between the first and last touches is denoted by $\Delta t_i^{\text{touch}} = t_{\text{last}} - t_{\text{first}}$. Calendar drift over this span is aggregated with rough drift, as specified below. After sampling ζ_{last} , the endpoint is

$$\zeta_{i+1} = \mathcal{D}_{t_{i+1}-t_{\text{last}}}(\zeta_{\text{last}}). \quad (\text{F.5})$$

Consequently calendar drift is applied separately only on the two no-touch portions. Its contribution over the middle span is counted once, inside the aggregated transition. If $\Delta \mathfrak{r}_i = 0$, the implementation uses the deterministic full-step flow directly.

F.3 Aggregated reciprocal transition

For $\Delta \mathfrak{r}_i > 0$ replace calendar time on the middle span by its average rate relative to rough time. The effective reciprocal parameters are

$$\vartheta_i^{\text{eff}} = \vartheta + \frac{2\alpha\Delta t_i^{\text{touch}}}{\boldsymbol{\Psi}_\zeta\Delta \mathfrak{r}_i}, \quad \mathfrak{J}_i^{\text{eff}} = \mathfrak{J} + \frac{\beta\Delta t_i^{\text{touch}}}{\Delta \mathfrak{r}_i}. \quad (\text{F.6})$$

Over $\Delta \mathfrak{r}_i$ the approximation is the ordinary affine diffusion with drift $\frac{1}{2}\boldsymbol{\Psi}_\zeta\vartheta_i^{\text{eff}} - \mathfrak{J}_i^{\text{eff}}\zeta$ and quadratic-variation rate $\boldsymbol{\Psi}_\zeta\zeta$. It has an exact Poisson–Gamma endpoint sampler. Define its noncentral contribution and Gamma scale by

$$r_{\zeta,i} = \zeta_{\text{first}} \exp\{-\mathfrak{J}\Delta \mathfrak{r}_i - \beta\Delta t_i^{\text{touch}}\}, \quad (\text{F.7})$$

$$c_{\zeta,i} = \frac{\boldsymbol{\Psi}_\zeta\Delta \mathfrak{r}_i}{2} \frac{1 - \exp\{-\mathfrak{J}\Delta \mathfrak{r}_i - \beta\Delta t_i^{\text{touch}}\}}{\mathfrak{J}\Delta \mathfrak{r}_i + \beta\Delta t_i^{\text{touch}}}. \quad (\text{F.8})$$

The ratio in the second expression equals one when its denominator is zero. The actual draw is

$$N_{\zeta,i} \sim \text{Poisson}\left(\frac{r_{\zeta,i}}{c_{\zeta,i}}\right), \quad \zeta_{\text{last}} = c_{\zeta,i} G_{\zeta,i}, \quad G_{\zeta,i} \mid N_{\zeta,i} \sim \text{Gamma}(\vartheta_i^{\text{eff}} + N_{\zeta,i}, 1). \quad (\text{F.9})$$

In particular,

$$\mathbb{E}_d[\zeta_{\text{last}} \mid \zeta_{\text{first}}, \text{clock information}] = r_{\zeta,i} + \vartheta_i^{\text{eff}} c_{\zeta,i}. \quad (\text{F.10})$$

The integrated rate in (F.8) avoids forming a very large effective rate and then multiplying it by a very small clock increment. The zero-clock case is handled before any division by $\Delta \mathfrak{r}_i$.

The assumptions $\alpha \geq 0$ and (4.73) imply

$$\vartheta_i^{\text{eff}} > 1, \quad \vartheta_{i,f}^{\text{eff}} := \vartheta_i^{\text{eff}} + \frac{2\mathfrak{N}}{\sqrt{\psi_\zeta}} > 1. \quad (\text{F.11})$$

Thus both numeraire characteristics of each aggregated substep remain in the intended interior regime. Exact sampling of this substep does not make the replacement of the two clocks by their average ratio exact.

F.4 Reciprocal conditional-return factor

First consider a smooth reciprocal transition over Δt , starting at ζ_0 and ending at ζ_1 . Its exact conditional return is $R(\zeta_1) = \mathbb{E}_d[S_1 \mid \zeta_1]/S_0$. We approximate this function by a normalized exponential tilt. The normalizer is evaluated under the endpoint law actually used in the sampler, including the Poisson mixture. This gives a positive approximation and allows both

$$\mathbb{E}_d[\widehat{R}] = 1, \quad \mathbb{E}_d[\zeta_1 \widehat{R}] = \mathbb{E}_f[\zeta_1] \quad (\text{F.12})$$

to hold exactly for that transition.

For the aggregated rough substep the domestic log moment-generating function is

$$K_{\zeta,i}(a) = -\vartheta_i^{\text{eff}} \log(1 - ac_{\zeta,i}) + \frac{ar_{\zeta,i}}{1 - ac_{\zeta,i}}, \quad ac_{\zeta,i} < 1. \quad (\text{F.13})$$

Its foreign target mean is

$$\bar{\zeta}_{i,f} = r_{\zeta,i} + \vartheta_{i,f}^{\text{eff}} c_{\zeta,i}. \quad (\text{F.14})$$

Choose $a_{\zeta,i}$ so that $K'_{\zeta,i}(a_{\zeta,i}) = \bar{\zeta}_{i,f}$. A form that avoids subtracting nearly equal quantities is

$$a_{\zeta,i} = \frac{2(\vartheta_{i,f}^{\text{eff}} - \vartheta_i^{\text{eff}})}{2\bar{\zeta}_{i,f} - \vartheta_i^{\text{eff}} c_{\zeta,i} + \sqrt{(\vartheta_i^{\text{eff}} c_{\zeta,i})^2 + 4r_{\zeta,i} \bar{\zeta}_{i,f}}}. \quad (\text{F.15})$$

Then the reciprocal return factor is

$$\log \widehat{R}_{\zeta,i} = a_{\zeta,i} \zeta_{\text{last}} - K_{\zeta,i}(a_{\zeta,i}). \quad (\text{F.16})$$

For $\mathfrak{N} = 0$ it equals one. For a zero rough increment it is also set to one. Because the post-touch flow is affine in ζ_{last} , the same two identities imply

$$\mathbb{E}_d[\zeta_{i+1} \widehat{R}_{\zeta,i} \mid \zeta_{\text{first}}, \text{clock information}] = \mathcal{D}_{t_{i+1}-t_{\text{last}}}(\bar{\zeta}_{i,f}). \quad (\text{F.17})$$

This is the foreign first moment of the *aggregated reciprocal substep*. It is not an assertion that the complete simulated foreign clock law is exact.

F.5 Residual reciprocal variance from the fair-variance formula

For the smooth reciprocal model, orthogonalizing the two Brownian drivers gives the pathwise decomposition

$$\begin{aligned} \log \frac{S_1}{S_0} &= \frac{\mathfrak{N}}{\sqrt{\Psi_\zeta}} \left(\log \frac{\zeta_1}{\zeta_0} + \mathfrak{N} \Delta t \right) - \frac{1 + \mathfrak{N} \sqrt{\Psi_\zeta} (\vartheta - 1)}{2} \int_0^{\Delta t} \frac{ds}{\zeta_s} \\ &\quad + \sqrt{1 - \mathfrak{N}^2} \int_0^{\Delta t} \frac{dW_s^\perp}{\sqrt{\zeta_s}}. \end{aligned} \quad (\text{F.18})$$

Conditional on the full ζ path, the last term is Gaussian with variance $(1 - \mathfrak{N}^2) \int_0^{\Delta t} ds/\zeta_s$. The simulation replaces this path integral by its domestic forward conditional expectation, in the same manner as the Heston forward-area approximation. This is a log-return variance, rather than a variance of the stock level.

For a smooth reciprocal process the fair variance is

$$\mathbb{E}_d \left[\int_0^{\Delta t} \frac{ds}{\zeta_s} \middle| \zeta_0 \right] = y \left(\frac{e^{\mathfrak{N} \Delta t} - 1}{\mathfrak{N} \zeta_0}; \vartheta, \Psi_\zeta \right), \quad (\text{F.19})$$

with argument $\Delta t/\zeta_0$ for $\mathfrak{N} = 0$. The function is characterized by

$$\frac{\Psi_\zeta}{2} x^2 y''(x) - \left[1 + \Psi_\zeta \left(\frac{\vartheta}{2} - 1 \right) x \right] y'(x) + 1 = 0, \quad y(0) = 0, \quad y'(0) = 1. \quad (\text{F.20})$$

An expression using the upper incomplete Gamma function is

$$y(x; \vartheta, \Psi_\zeta) = \frac{2}{\Psi_\zeta} \int_0^1 u^{\vartheta-2} \Gamma \left(0, \frac{2(1-u)}{\Psi_\zeta x} \right) du, \quad x > 0, \quad \vartheta > 1. \quad (\text{F.21})$$

All terms are nonnegative, and the logarithmic singularity at $u = 1$ is integrable. Equivalently,

$$y'(x) = \frac{2}{\Psi_\zeta x} \int_0^1 u^{\vartheta-2} \exp \left\{ -\frac{2(1-u)}{\Psi_\zeta x} \right\} du. \quad (\text{F.22})$$

These forms avoid a confluent hypergeometric evaluation. They also specify the continuous small-step limit $y(x) \sim x$.

For numerical evaluation the implementation uses the equivalent integration-by-parts form

$$y(x; \vartheta, \Psi_\zeta) = \frac{2}{\Psi_\zeta (\vartheta - 1)} \left[\Gamma \left(0, \frac{2}{\Psi_\zeta x} \right) + \int_0^1 \frac{1 - (1-u)^{\vartheta-1}}{u} \exp \left\{ -\frac{2u}{\Psi_\zeta x} \right\} du \right]. \quad (\text{F.23})$$

The subtraction in the numerator is computed with `expm1` and `log1p`. To resolve the narrow peak for small steps, the integration coordinate is rescaled to $2u/(\boldsymbol{\psi}_\zeta x)$ and integrated up to the smaller of $2/(\boldsymbol{\psi}_\zeta x)$ and 48, using 64-point Gauss–Legendre quadrature. The discarded portion is an exponentially small tail. The incomplete Gamma term is evaluated by its convergent series when its argument is at most one, and by Gauss–Laguerre quadrature otherwise. These are numerical approximations to (F.23); their accuracy should be checked over the parameter range used in calibration, particularly near $\vartheta = 1$.

For a touched step use the domestic effective parameters from (F.6) and set

$$\widehat{V}_{\zeta,i}^\perp = (1 - \mathfrak{n}^2) y \left(\frac{\exp\{\boldsymbol{\mathfrak{z}}_i^{\text{eff}} \Delta \mathfrak{r}_i\} - 1}{\boldsymbol{\mathfrak{z}}_i^{\text{eff}} \zeta_{\text{first}}}; \vartheta_i^{\text{eff}}, \boldsymbol{\psi}_\zeta \right). \quad (\text{F.24})$$

The argument is $\Delta \mathfrak{r}_i / \zeta_{\text{first}}$ when the effective rate is zero; the variance itself is zero when the rough increment is zero. The numerical evaluator must return a finite nonnegative value. Its approximation error changes dispersion but does not change the conditional mean of the compensated lognormal multiplier below.

This forward projection averages over the reciprocal path. It does not condition the integrated inverse state on the realised ζ_{last} . Nor does it retain the fluctuation of the integral in the second term of (F.18). In particular, when $|\mathfrak{m}| = 1$ the independent Gaussian disappears, but conditioning only on endpoints does not remove all true bridge randomness. The method is therefore a conditional-return and residual- variance approximation, not an exact reciprocal bridge sampler.

F.6 Heston return normalized under the sampled-touch law

The calendar Heston contribution uses the residual variance in (E.3). On a touched step, its forward-area projection restarts from zero at the last touch. Thus

$$\widehat{V}_{\xi,i}^\perp = \frac{(1 - \rho^2) \mathcal{H}}{2} A_{\xi,i}. \quad (\text{F.25})$$

The area formulas are those of the preceding implementation appendix; the post-touch duration is $t_{i+1} - t_{\text{last}}$. In particular this approximation discards the pre-last-touch area on a touched step.

The return normalizer, however, must use the law under which the endpoint was actually drawn. Conditional on the retained pending touch and the sampled last-touch information, but *before drawing its auxiliary Poisson count*, that law has the form

$$N_{\xi,i} \sim \text{Poisson}(\ell_{\xi,i}), \quad \xi_{i+1} \mid N_{\xi,i} \sim c_{\xi,i} \text{Gamma}(\nu_{\xi,i} + N_{\xi,i}, 1). \quad (\text{F.26})$$

For continuation, $\nu_{\xi,i} = 2 - \theta$ and the scale and Poisson mean come from the conditional touch-time transition. For a touched step, $\nu_{\xi,i} = 1$, $\ell_{\xi,i} = 0$, and the scale is that of the final meander. Consequently the correct conditional log MGF is

$$K_{\xi,i}(a) = -\nu_{\xi,i} \log(1 - ac_{\xi,i}) + \ell_{\xi,i} \frac{ac_{\xi,i}}{1 - ac_{\xi,i}}, \quad ac_{\xi,i} < 1. \quad (\text{F.27})$$

It is essential to integrate over the Poisson count in this normalizer. Conditioning on a realised count would normalize a different conditional approximation.

The Heston slope $a_{\xi,i}$ is inherited from the ordinary CIR exponential-tilt prescription: before conditioning on touch metadata, it would give the foreign endpoint mean

$$\mathbb{E}_f[\xi_{i+1} \mid \xi_i] = \xi_i e^{-(\kappa - \rho\sqrt{\mathcal{H}})\Delta t_i} + \theta \frac{1 - e^{-(\kappa - \rho\sqrt{\mathcal{H}})\Delta t_i}}{\kappa - \rho\sqrt{\mathcal{H}}}, \quad (\text{F.28})$$

with the continuous zero-rate limit. The new return factor is obtained as follows. For the ordinary, unconditioned CIR law define

$$c_{\xi,i}^{\text{CIR}} = \frac{1 - e^{-\kappa\Delta t_i}}{\kappa}, \quad r_{\xi,i}^{\text{CIR}} = \xi_i e^{-\kappa\Delta t_i}, \quad c_{\xi,i}^{\text{CIR}} = \Delta t_i \quad (\kappa = 0), \quad (\text{F.29})$$

and let $\bar{\xi}_{i,f}$ be the right-hand side of (F.28). Then

$$a_{\xi,i} = \frac{2(\bar{\xi}_{i,f} - r_{\xi,i}^{\text{CIR}} - \theta c_{\xi,i}^{\text{CIR}})}{c_{\xi,i}^{\text{CIR}} \left[2\bar{\xi}_{i,f} - \theta c_{\xi,i}^{\text{CIR}} + \sqrt{(\theta c_{\xi,i}^{\text{CIR}})^2 + 4r_{\xi,i}^{\text{CIR}} \bar{\xi}_{i,f}} \right]}. \quad (\text{F.30})$$

The slope is used with the *conditional* MGF (F.27), rather than the ordinary CIR normalizer:

$$\log \hat{R}_{\xi,i} = a_{\xi,i} \xi_{i+1} - K_{\xi,i}(a_{\xi,i}). \quad (\text{F.31})$$

This preserves the unit conditional mean for the sampled-touch law, including its dependence on a retained future touch. It does not also force (F.28) under each such conditional law. When $\rho = 0$ or $\mathcal{H} = 0$, the factor is one.

For either endpoint factor, $ac < 1$ is the domain of the first-moment normalizer. The stronger $2ac < 1$ ensures a finite second moment of that factor under its endpoint law. These are numerical restrictions on the exponential ansatz, distinct from the model's two Feller restrictions. The implementation rejects coarse steps failing $2ac < 1$ and requests grid refinement; it does not modify the slope. Neither a clipped endpoint return nor a sample-mean renormalization supplies the stated moment identities. Finiteness of a factor's second moment alone is not a general proof that every full-model payoff has finite variance.

F.7 Asset update and its martingale identity

After the two endpoints have been sampled, the complete step is

$$\begin{aligned} \log \frac{S_{i+1}}{S_i} &= \log \hat{R}_{\xi,i} + \log \hat{R}_{\zeta,i} - \frac{1}{2} \hat{V}_{\xi,i}^\perp - \frac{1}{2} \hat{V}_{\zeta,i}^\perp \\ &\quad + \sqrt{\hat{V}_{\xi,i}^\perp} Z_{\xi,i}^\perp + \sqrt{\hat{V}_{\zeta,i}^\perp} Z_{\zeta,i}^\perp, \end{aligned} \quad (\text{F.32})$$

where the two standard normals are independent of each other and of all state draws on that step. They remain conditionally independent given the sampled state and touch information. The same variance value is used in each Gaussian coefficient and its corresponding $-\frac{1}{2}V$ compensator.

To see the discrete martingale property, first average over the two normals. Both compensated exponentials have conditional mean one. Next average over the reciprocal Poisson–Gamma endpoint, holding the sampled first/last-touch geometry, edge increment and ξ_{i+1} fixed: by (F.13)–(F.16), $\mathbb{E}[\hat{R}_{\zeta,i} \mid \text{that information}] = 1$. On a touched step the new *future* pending touch is sampled from ξ_{i+1} , but it does not enter either return factor on the present step. Its conditional law can therefore be integrated out. Finally average over the ξ endpoint conditional on the *entering* pending touch and sampled last-touch geometry, and then over that last-touch geometry. Equation (F.27) gives a unit mean for $\hat{R}_{\xi,i}$ in this final endpoint average. Therefore

$$\boxed{\mathbb{E}[S_{i+1} \mid \mathcal{F}_i^{\text{sim}}] = S_i, \quad \mathbb{E}[S_{t_j}] = S_0,} \quad (\text{F.33})$$

where $\mathcal{F}_i^{\text{sim}}$ includes the simulator state and its retained pending touch. This is an analytic identity for the numerical transition, up to evaluation and floating-point error. It does not require the Monte Carlo mean of any realised batch to equal S_0 .

Normalizing relative to the retained touch metadata is stronger than normalizing only an unconditional CIR transition. It changes the approximation to the joint asset–clock law: it does not construct the foreign-measure distribution of future touch times. Thus (F.33) establishes the desired discrete martingale property, but alone establishes neither exact Heston prices nor convergence of the complete joint scheme to (4.64).

F.8 Streaming, calibration example, and diagnostics

In order, one iteration performs the following operations:

1. Sample the continuing or touched clock transition, retaining the conditional endpoint-law parameters needed for $K_{\xi,i}$ before the auxiliary Poisson draw.

2. Propagate ζ deterministically on a no-touch step. Otherwise apply the pre-touch flow, sample the aggregated reciprocal transition, and apply the post-touch flow.
3. Compute the normalized Heston and reciprocal return factors and their non-negative forward residual variances.
4. Apply (F.32). At an expiry, evaluate all requested strikes and call/put signs from the same stock sample, together with its unadjusted forward estimate.

Only current states and requested expiry outputs are retained. No full stock, reciprocal, or clock trajectory is needed for pricing.

The two residual Gaussians use four sign combinations for each base path: $(+, +)$, $(+, -)$, $(-, +)$, and $(-, -)$. All four members share their touch, ξ , and ζ draws. Independence holds between base-path groups, not between members of a group. Monte Carlo standard errors are therefore computed from averages of the four payoffs within each group.

The example notebook is

`py/example_rough_reciprocal_market_calibration.ipynb`.

The complete parameter descriptor follows setup at the top of the notebook; initial values are obtained by evaluating that descriptor. The notebook displays the default-guess smiles before calibration. As in the Heston example, the saved best and the defaults are repriced with the current settings; the lower finite objective supplies the starting point. If that selected point differs from the displayed guess, its smiles are shown before the first optimizer probe. Every improved point is saved atomically, and an interrupt returns the best completed result. Rerunning the fit cell resumes from the saved physical parameters with a fresh optimizer. Result plots and diagnostics can also read that checkpoint after a kernel restart without recalibration. The live display retains the latest and best errors. The default budget is 512 residual evaluations, including derivative probes, with the same stopping tolerances as the Heston example; reaching a numerical stopping condition does not establish fit quality. The subsequent comparison uses the calibrated parameters and reports raw simulated forwards and their group-based standard errors, in addition to option-price or implied-volatility residuals. A forward diagnostic uses

$$\frac{\widehat{\mathbb{E}}[S_T] - S_0}{\text{SE}(\widehat{\mathbb{E}}[S_T])}, \quad (\text{F.34})$$

without modifying paths or payoffs to remove the numerator. Finite sample fluctuations should be distinguished from systematic error across seeds and path counts.

Focused checks for this implementation cover copied clock paths, deterministic no-touch drift and its zero-rate limits, aggregated Poisson–Gamma endpoint moments, both reciprocal tilt identities, sampled-touch Heston normalizers, and the fair-variance evaluator. Monte Carlo diagnostics then exercise the complete update, including both signs of the correlations and steps with and without touches. Refining the calendar grid probes the drift aggregation, elapsed-edge-time and bridge approximations. Increasing path count probes sampling error. These are separate checks: passing the martingale diagnostic does not remove the other approximation errors described above.

G Calendar-time asset simulation with a rough reciprocal state

This appendix describes the separate implementation of (4.65), with $c_c = 1$ and $\Psi_S = 0$. The sampled touches, ξ transitions, and elapsed-edge-time approximation are retained from Appendix F, including (F.1). The reciprocal endpoint uses the same exact calendar flows outside the touched span and the aggregated rough affine transition inside it. Thus calendar drift is included once: explicitly on each outer interval and through the effective coefficients on the middle interval. These shared constructions are not repeated here.

Calendar reciprocal variance. The asset now requires $\int dt/\zeta$, including during flat-clock intervals. For a calendar-only interval of length d , beginning at $\zeta > 0$, its value is obtained exactly from the deterministic flow:

$$I(\zeta, d) = \int_0^d \frac{ds}{\mathcal{D}_s(\zeta)} = \begin{cases} \frac{1}{\alpha} \log \left(1 + \frac{\alpha(e^{\beta d} - 1)}{\beta \zeta} \right), & \alpha > 0, \beta \neq 0, \\ \frac{1}{\alpha} \log \left(1 + \frac{\alpha d}{\zeta} \right), & \alpha > 0, \beta = 0, \\ \frac{e^{\beta d} - 1}{\beta \zeta}, & \alpha = 0, \beta \neq 0, \\ \frac{d}{\zeta}, & \alpha = \beta = 0. \end{cases} \quad (\text{G.1})$$

Continuous small-rate limits are evaluated with numerically stable exponential and logarithmic functions. For a touched step, denote the reciprocal states immediately before and after the middle transition by ζ_{first} and ζ_{last} . For $\Psi_\zeta > 0$ the calendar inverse-state area on that span is approximated by a positive trapezoid:

$$\begin{aligned} \widehat{I}_{\zeta,i} &= I(\zeta_i, t_{\text{first}} - t_i) \\ &\quad + \frac{t_{\text{last}} - t_{\text{first}}}{2} \left(\frac{1}{\zeta_{\text{first}}} + \frac{1}{\zeta_{\text{last}}} \right) \\ &\quad + I(\zeta_{\text{last}}, t_{i+1} - t_{\text{last}}). \end{aligned} \quad (\text{G.2})$$

When $\Delta \mathfrak{r}_i = 0$, the whole step instead uses the unsplit exact calendar flow and $I(\zeta_i, \Delta t_i)$. This also covers a touched span with zero elapsed edge time. The trapezoid approximates a calendar-time bridge integral; it is not the conditional expectation of that integral. Grid refinement is therefore needed even when the endpoint law and the no-touch integrals are accurate.

Asset update and conditional pricing. Let $\widehat{R}_{\xi,i}$ be the analytically normalized Heston return factor for the actual retained-touch transition, as in Appendix F. Let $\widehat{I}_{\xi,i}$ be the Heston dispersion-area approximation used there. The calendar-stock update has just one return tilt:

$$\begin{aligned}\widehat{V}_i &= (1 - \rho^2) \frac{\mathcal{H}}{2} \widehat{I}_{\xi,i} + \widehat{I}_{\zeta,i}, \\ \frac{S_{i+1}}{S_i} &= \widehat{R}_{\xi,i} \exp \left(-\frac{\widehat{V}_i}{2} + \sqrt{\widehat{V}_i} Z_i \right), \quad Z_i \sim N(0, 1).\end{aligned}\tag{G.3}$$

Here Z_i is independent of the sampled state and touch information. There is no reciprocal-state return tilt and no $\mathfrak{D}$ parameter. Conditional on that information, the exponential has mean one. The analytic normalization of $\widehat{R}_{\xi,i}$ then gives conditional mean S_i before the step is sampled, within the normalizer's finite moment domain. This is a property of the discrete approximation; it does not replace a martingale argument for the continuous hybrid model. No realized sample-mean correction is applied.

For European option pricing the independent Gaussian residuals can be integrated analytically. Each path retains the sum of the log-return tilts and the sum of the residual variances; a conditional lognormal option formula is evaluated at each expiry. Thus the residual normals need not be sampled or replicated with antithetic signs. Forward diagnostics use the unadjusted mean of the resulting conditional forwards.

Implementation and forward-variance checks. The separate modules in `py/` are:

- `rough_vol_reciprocal.py`: state and calendar integral;
- `rough_vol_reciprocal_pricing.py`: conditional option pricing;
- `rough_vol_reciprocal_market.py`: calibration;
- `rough_vol_reciprocal_vix.py`: conditional forward variance.

The example notebook is

`py/example_rough_vol_reciprocal_market_calibration.ipynb`.

Its model identifier and saved calibration results are separate from those of the rough-asset implementation.

The implemented state descriptor uses finite ϑ and $\mathfrak{V}_{\zeta} \geq 0$. Setting $\mathfrak{V}_{\zeta} = 0$ in this descriptor also removes the constant rough drift $\mathfrak{V}_{\zeta}\vartheta/2$; the remaining

aggregated middle flow and its inverse integral are deterministic conditional on the clock. The positive- $\aleph$ drift-only model (4.79) requires an independent rough intercept and is not obtained by that single parameter setting. The general positive- Ψ_S hybrid asset also requires the additional rough return channel; the calendar implementation does not simulate it.

A further separate implementation, the `rough_vol_reciprocal_drift` family of modules, instead fits the independent intercept $\aleph$. On the middle interval its aggregated source is $\alpha(t_{\text{last}} - t_{\text{first}}) + \aleph \Delta \mathfrak{r}_i$. At $\Psi_\zeta = 0$ this source is retained in the exact aggregated deterministic flow and inverse-state integral. Randomness remains through the sampled rough clock. For positive diffusion, the calendar inverse-area approximation is the exact deterministic aggregated area multiplied by

$$\frac{1/\zeta_{\text{first}} + 1/\zeta_{\text{last}}}{1/\zeta_{\text{first}} + 1/\zeta_{\text{last}}^{\text{det}}},$$

where $\zeta_{\text{last}}^{\text{det}}$ is the endpoint under that same aggregated drift with diffusion suppressed. This positive correction recovers the deterministic integral continuously as diffusion vanishes; it remains an endpoint approximation for a stochastic bridge and requires grid refinement. The fitter uses separate dimensionless coordinates for $(\aleph - \Psi_\zeta/2)U/\zeta_0$ and $\Psi_\zeta U/\zeta_0$, where U is the chosen deterministic reference unit of rough time, not its expected realized increment. This enforces the strict positive-diffusion constraint without a Feller-ratio coordinate, while permitting exact zero-diffusion profiles. Its results notebook is

`py/example_rough_vol_reciprocal_drift_joint_calibration.ipynb`.

Forward variance can be checked by simulating fresh future clock and reciprocal paths conditional on the current observed (ξ_t, ζ_t) , using the appropriate calendar origin, and averaging their calendar inverse-state areas. The Heston area is evaluated analytically as in (4.81). A pending future touch retained internally by an earlier sampler is not part of the observed conditioning state. Comparing this estimate with the frozen- ζ approximation measures the effect of future reciprocal dynamics. Independent seeds and larger path counts assess sampling error; smaller steps assess the clock, drift-aggregation, and inverse-area approximations. Ordinary standard error estimates require finite second moments: for the stochastic rough affine transitions, $\vartheta > 2$ is a sufficient interior choice for the endpoint inverse-square moment.

First-touch lower bound for VIX. Retain $c_c = 1$, $\Psi_S = 0$, so the stock has no rough-clock quadratic variation. Let $T_t = \inf\{s \geq t : \xi_s = 0\}$ be the next

boundary contact, conditional on the current observed state. Before T_t the clock $\mathfrak{r}$ is flat and the reciprocal state follows its calendar flow:

$$\mathcal{D}_u(\zeta_t) = \begin{cases} \zeta_t e^{-\beta u} + \frac{\alpha}{\beta}(1 - e^{-\beta u}), & \beta \neq 0, \\ \zeta_t + \alpha u, & \beta = 0. \end{cases} \quad (\text{G.4})$$

For $0 < \theta < 1$ and $u > 0$, its exact first-touch survival probability is

$$\mathbb{P}_t(T_t > t + u) = \begin{cases} \frac{\gamma\left(1 - \theta, \frac{\kappa \xi_t}{e^{\kappa u} - 1}\right)}{\Gamma(1 - \theta)}, & \kappa \neq 0, \\ \frac{\gamma(1 - \theta, \xi_t/u)}{\Gamma(1 - \theta)}, & \kappa = 0, \end{cases} \quad (\text{G.5})$$

where γ is the lower incomplete Gamma function. Dropping only the nonnegative reciprocal area after T_t , while retaining the full exact conditional Heston area, gives

$$\begin{aligned} \frac{\text{VIX}_t^2}{100^2} \geq \frac{1}{\Delta t} \left\{ \frac{\mathcal{H}}{2} \left[\xi_t \frac{1 - e^{-\kappa \Delta t}}{\kappa} + \frac{\theta}{\kappa} \left(\Delta t - \frac{1 - e^{-\kappa \Delta t}}{\kappa} \right) \right] \right. \\ \left. + \int_0^{\Delta t} \frac{\mathbb{P}_t(T_t > t + u)}{\mathcal{D}_u(\zeta_t)} du \right\}. \end{aligned} \quad (\text{G.6})$$

The Heston area in square brackets has limit $\xi_t \Delta t + \theta \Delta t^2/2$ at $\kappa = 0$. The inequality requires only positivity of ζ after contact, not a particular sign of its rough drift or independence of the states.

The integral is an analytic lower bound; its finite-node quadrature is not a certified bound without a numerical error estimate. Its square root can be used as a VIX approximation only when the omitted post-touch reciprocal area is negligible, as may occur after a fast upward reset of ζ . Accuracy is not automatic when rough dynamics are absent or when rough activity instead lowers ζ . Nested conditional simulation should check this omission across states and expiries, with separate sampling-error and grid-refinement diagnostics. Rescaling simulated VIX levels to a parity-implied forward does not preserve the lower-bound ordering and does not remove state-dependent approximation error.

H Vol-of-Vol Term-Structure




Fig 2: *Smooth* contribution ($\xi = 0.01$)




Fig 3: *Rough* contribution ($\xi = 0.01$)




Fig 4: *Smooth* contribution ($\xi = 0.1$)




Fig 5: *Rough* contribution ($\xi = 0.1$)




Fig 6: *Smooth* contribution ($\xi = 0.2$)




Fig 7: *Rough* contribution ($\xi = 0.2$)




Fig 8: *Smooth* contribution ($\xi = 0.3$)




Fig 9: *Rough* contribution ($\xi = 0.3$)




Fig 10: *Smooth* contribution ($\xi = 0.5$)




Fig 11: *Rough* contribution ($\xi = 0.5$)




Fig 12: *Smooth* vol-of-vol ($\xi = 1$)




Fig 13: *Rough* vol-of-vol ($\xi = 2$)




Fig 14: *Smooth* vol-of-vol ($\xi = 1$)




Fig 15: *Rough* vol-of-vol ($\xi = 2$)

I Simulation Results

We ran simulation for the 3 years period with 2^9 time steps and 2^{19} paths for the following set of model parameters

$$\aleph^2 = \frac{1}{52}, \quad \nu = 0.03, \quad \rho = -0.99, \quad \alpha = 1, \quad \beta = 0, \quad \gamma = 0.3, \quad \xi_0 = 0.5\aleph^2$$

In the following, we plot the resulting implied volatilities against $\log \frac{K}{F}$. The plots' width is 2 standard deviations computed with ATMF volatilities.

References

- [1] Lorenzo Bergomi. *Volatility modeling: state of the art & outlook*. ENSAE, Paris, January 29, 2020




Fig 16: Implied Vol: $T = 2$ days




Fig 17: Implied Vol: $T = 1$ weeks




Fig 18: Implied Vol: $T = 2$ weeks




Fig 19: Implied Vol: $T = 1$ month

- [2] Jim Gatheral, Thibault Jaisson, and Mathieu Rosenbaum. *Volatility is rough*. Quantitative Finance 18 (6): 933–949, 2018. <https://doi.org/10.1080/14697688.2017.1393551>.
- [3] Mikkel Bennedsen, Asger Lunde, and Mikko S. Pakkanen. *Decoupling the Short- and Long-Term Behavior of Stochastic Volatility*. Journal of Financial Econometrics 20 (5): 961–1006, 2022. <https://doi.org/10.1093/jjfinec/nbaa049>.
- [4] Sigurd Emil Rømer. *Empirical analysis of rough and classical stochastic volatility models to the SPX and VIX markets*. Quantitative Finance 22 (10): 1805–1838, 2022. <https://doi.org/10.1080/14697688.2022.2081592>.




Fig 20: Implied Vol: $T = 2$ months




Fig 21: Implied Vol: $T = 3$ months




Fig 22: Implied Vol: $T = 6$ months




Fig 23: Implied Vol: $T = 9$ months

- [5] Giuseppe Brandi and Tiziana Di Matteo. *Multiscaling and rough volatility: An empirical investigation*. International Review of Financial Analysis 84: 102324, 2022. <https://doi.org/10.1016/j.irfa.2022.102324>.
- [6] Carsten H. Chong, Thomas Delerue, and Guoying Li. *When Frictions are Fractional: Rough Noise in High-Frequency Data*. Journal of the American Statistical Association 120 (551): 1531–1544, 2025. <https://doi.org/10.1080/01621459.2024.2428466>.
- [7] Gregory Pelts. *Unspanned Stochastic Volatility, Conformal Symmetries, and Stochastic Time*. <https://ssrn.com/abstract=3001967>, June 29, 2017
- [8] Robert C. Merton, *On the Pricing of Corporate Debt: The Risk Structure of Interest Rates*. Journal of Finance. 29 (2): 449–470, 1974




Fig 24: Implied Vol: $T = 1$ year




Fig 25: Implied Vol: $T = 1\frac{1}{2}$ year




Fig 26: Implied Vol: $T = 2$ year




Fig 27: Implied Vol: $T = 3$ year

- [9] R. N. Pillai. *On Mittag-Leffler functions and related distributions*. Annals of the Institute of Statistical Mathematics 42 (1): 157–161, 1990.
- [10] R. Gorenflo, A. A. Kilbas, F. Mainardi, S. V. Rogosin. *Mittag-Leffler Functions, Related Topics and Applications*. Springer Monographs in Mathematics, 2nd ed., Springer, Berlin, 2020.
- [11] N. H. Bingham. *Limit theorems for occupation times of Markov processes*. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete 17: 1–22, 1971.
- [12] Mark M. Meerschaert and Peter Straka. *Inverse stable subordinators*. Mathematical Modelling of Natural Phenomena 8 (2): 1–16, 2013.

- [13] V. M. Zolotarev. *One-Dimensional Stable Distributions*. Translations of Mathematical Monographs, Vol. 65, American Mathematical Society, Providence, RI, 1986.
- [14] Carole Bernard, Zhenyu Cui, and Don McLeish. *On the martingale property in stochastic volatility models based on time-homogeneous diffusions*. 2014. <https://arxiv.org/abs/1310.0092>.
- [15] Peter Carr, Travis Fisher, and Johannes Ruf. *On the Hedging of Options On Exploding Exchange Rates*. 2013. <https://arxiv.org/abs/1202.6188>.